

Essays in Macroeconomics and Finance

Kyle Jurado

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2015

© 2015

Kyle Jurado

All rights reserved

ABSTRACT

Essays in Macroeconomics and Finance

Kyle Jurado

This dissertation contains three essays in macroeconomics and finance. Chapter 1 estimates the relative importance of agents receiving advance information or having distorted beliefs about future fundamentals in explaining a set of macroeconomic and financial data. Chapter 2 proposes a new measure of time-varying aggregate uncertainty, which is based on information from a large panel of macroeconomic and financial time series. Chapter 3 explores how imposing risk constraints on financial intermediaries in a continuous-time heterogeneous-agent economy can affect equilibrium allocations and asset price dynamics.

Fluctuations in the beliefs of economic agents can be driven by current fundamentals, advance information about future fundamentals, or distortions resulting from informational or psychological limitations. Chapter 1 presents a dynamic stochastic general equilibrium (DSGE) model that jointly considers all three possibilities and estimates their relative importance for explaining macroeconomic and financial data. In the model, agents' beliefs are based on a subjective probability measure that can be stochastically distorted relative to the objective measure. By directly parameterizing the Radon-Nikodým process linking these two measures, it is possible to solve the model under the subjective measure and change back to the objective measure for estimation. To help the model jointly explain macroeconomic and financial data, it features recursive preferences alongside a number of more standard business-cycle frictions. To facilitate estimation, a solution method is presented that allows risk premia to affect the model's linear-approximate dynamics. To discipline the estimates, direct data on subjective forecasts is included in the set of observable variables. Results indicate that both advance information and distorted beliefs are important. On average about two-thirds of the fluctuations in endogenous variables can be attributed to these two sources. While they are equally important for explaining output and employment, advance information is most important for explaining inflation and investment, and distorted beliefs

are most important for explaining stock returns and consumption.

Chapter 2 exploits a data rich environment to provide direct econometric estimates of time-varying macroeconomic uncertainty. The estimates display significant independent variations from popular uncertainty proxies, suggesting that much of the variation in the proxies is not driven by uncertainty. Quantitatively important uncertainty episodes appear far more infrequently than indicated by popular uncertainty proxies, but when they do occur, they are larger, more persistent, and more correlated with real activity. The estimates provide a benchmark to evaluate theories for which uncertainty shocks play a role in business cycles.

Finally, Chapter 3 studies the general equilibrium effects of introducing a Value-at-Risk (VaR) constraint into a dynamic continuous-time economy with homogeneous preferences, inefficient endogenous volatility, fire sales, and economically valuable financial intermediation. The main finding is that through its impact on the stationary distribution of wealth in the economy, a VaR constraint can reduce the average level of endogenous volatility and lower the probability of entering a crisis regime. It does so by forcing agents to sell off their risky asset holdings earlier than they otherwise would, while they still have a large equity buffer to absorb losses. This chapter is the first study to explore the effects of a VaR constraint in a model that does not feature any heterogeneity in preferences or beliefs, and in which endogenous volatility and crises are socially inefficient.

Contents

List of Figures	vii
List of Tables	viii
Acknowledgments	x
1 Advance Information and Distorted Beliefs	1
1.1 Introduction	1
1.2 Related Literature	5
1.3 Sources of Fluctuations in Beliefs	10
1.3.1 A Permanent-Income Example	10
1.3.2 Current Fundamentals	12
1.3.3 Advance Information	13
1.3.4 Distorted Beliefs	14
1.3.5 Combined System	17
1.3.6 Identification: Key Moments	18
1.3.7 Correlations and Measurement Errors	21
1.3.8 Two Monte Carlo Experiments	23
1.4 Quantitative Equilibrium Model	25
1.4.1 Household	25
1.4.2 Labor Union	28
1.4.3 Productive Firm	30
1.4.4 Retail Firm	33

1.4.5	Government	34
1.4.6	Market Clearing and Equilibrium	35
1.4.7	Fundamental Processes	36
1.5	Solution Methods	37
1.5.1	Solving Models with Distorted Beliefs	38
1.5.2	Linear Approximations and Recursive Preferences	41
1.6	Empirical Analysis	45
1.6.1	Estimation Method	45
1.6.2	Data	47
1.6.3	Calibrated Parameters and Priors	48
1.6.4	Estimation Results	52
1.6.5	Decomposing Fluctuations	57
1.7	Robustness	62
1.7.1	No Stock Returns	62
1.7.2	Measurement Errors in Survey Forecasts	64
1.7.3	No Distorted Beliefs	65
1.7.4	More Nominal Rigidity	66
1.8	Two Historical Episodes	67
1.8.1	Construction of Historical Decompositions	68
1.8.2	The Dot-Com Boom and Bust	68
1.8.3	The Great Recession	72
1.9	Conclusion	75
2	Measuring Uncertainty	77
2.1	Introduction	77
2.2	Related Empirical Literature	84
2.3	Econometric Framework	87
2.3.1	Construction of Forecast Uncertainty	89

2.3.2	Time-varying Uncertainty: A Statistical Decomposition	92
2.4	Empirical Implementation and Macro Data	95
2.5	Estimates of Macro Uncertainty	99
2.5.1	The Role of the Predictors	105
2.5.2	Uncertainty Versus Stock Market Volatility	110
2.5.3	Macro Uncertainty and Macroeconomic Dynamics	113
2.5.4	Comparison with Measures of Dispersion	122
2.6	Results: Firm-Level Common Uncertainty	125
2.7	Conclusion	129
3	Volatility, Crises, and Value-at-Risk	132
3.1	Introduction	132
3.2	Unconstrained Model	134
3.2.1	Environment	135
3.2.2	Pareto optimality	139
3.2.3	Competitive equilibrium	140
3.2.4	Numerical example	147
3.3	Adding Value-at Risk	148
3.3.1	Constrained equilibrium	150
3.3.2	Numerical example	154
3.4	Stationary Distribution	157
3.4.1	Unconstrained model	158
3.4.2	Constrained model	159
3.4.3	Robustness and limitations	163
3.5	Conclusion	165
	Bibliography	166

A	Appendix for Chapter 1	182
A.1	Data Description	182
A.1.1	Raw Data	182
A.1.2	Data Transformations	184
A.2	Related Types of Disturbances	185
A.2.1	Ambiguity	185
A.2.2	Confidence	186
A.2.3	Noise Traders	187
A.3	Belief Distortions from Noisy Signals	188
A.3.1	Signals about the Future	188
A.3.2	Limited Information	190
A.4	An Optimal Policy Example	193
B	Appendix for Chapter 2	197
B.1	Robustness	197
B.1.1	Macro Uncertainty Factor	197
B.1.2	Alternative Estimates of Uncertainty	199
B.1.3	Recursive Out-of-Sample Estimation	203
B.1.4	Bloom (2009) VAR	207
B.2	Data Appendix	210
B.2.1	Macro Dataset	211
B.2.2	Financial Dataset	217
B.2.3	Firm-level Dataset	222
B.2.4	Data for VAR Analysis	223
C	Appendix for Chapter 3	226
C.1	Proof of Proposition 1	226
C.2	Proof of Proposition 2	228

C.3	Proof of Proposition 3	229
C.4	Proof of Proposition 4	231
C.5	Proof of Proposition 5	232
C.6	Proof of Proposition 6	235
C.7	Numerical methods	236

List of Figures

1.1	Monte Carlo experiments	24
1.2	Dot-Com boom and bust (1992-2002)	70
1.3	Run-up to the Great Recession (2002-2009)	73
2.1	Aggregate uncertainty: $\overline{\mathcal{U}}_t^y(h)$ for $h = 1, 3, 12$	100
2.2	Predictor uncertainty	106
2.3	The role of predictors	108
2.4	Uncertainty in the S&P 500 index	110
2.5	Stock market implied volatility and uncertainty	112
2.6	Impulse responses from an eleven-variable VAR(12)	117
2.7	Impulse responses from an eight-variable VAR(12)	118
2.8	Cross-sectional dispersion and uncertainty	124
2.9	Impulse responses using dispersion-based proxies	126
2.10	Firm-level uncertainty	128
3.1	Competitive equilibrium	149
3.2	Constrained equilibrium	155
3.3	Stationary density $p(\eta)$ in the unconstrained economy	160
3.4	Stationary density $p(\eta)$ in the constrained economy	162
B.1	Different estimates of macro uncertainty when $h = 1$	201
B.2	Percentile-based estimates of aggregate uncertainty when $h = 1$	202

B.3	EGARCH aggregate uncertainty	204
B.4	Uncertainty factor based on recursive forecasts	206
B.5	Eight-variable VAR(12) with uncertainty ordered last	208

List of Tables

1.1	Data	47
1.2	Calibrated parameters	49
1.3	Prior distributions	51
1.4	Estimated structural parameters	52
1.5	Estimated persistence parameters of exogenous processes	54
1.6	Estimated volatility parameters of exogenous processes	55
1.7	Selected moments in the model and data	56
1.8	Variance decomposition	58
1.9	Variance decomposition for selected exogenous processes	60
1.10	Variance decomposition at the prior mean	61
1.11	Variance decomposition without stock returns	63
1.12	Variance decomposition with classical measurement error in survey forecasts	64
1.13	Variance decomposition without distorted beliefs	66
1.14	Variance decomposition with more nominal rigidity	67
2.1	Summary statistics	101
2.2	Cross-sectional averages of R^2	104
2.3	Decomposition of variance	120
2.4	Cross-sectional averages of R^2	130
3.1	Baseline parameter configuration	163

3.2 Sensitivity analysis	164
B.1 Different aggregation methods	200
B.2 Variance decompositions with uncertainty ordered last	209
B.3 Output and income	212
B.4 Labor market	213
B.5 Housing	214
B.6 Consumption, orders, and inventories	214
B.7 Money and credit	214
B.8 Bond and exchange rates	215
B.9 Prices	216
B.10 Stock market	216
B.11 Financial data	218

Acknowledgments

With grateful acknowledgement to my dissertation sponsor, Ricardo Reis. His guidance has been invaluable to me. I thank Michael Woodford and Serena Ng for their support since my time as an undergraduate student at Columbia. They inspired me to pursue a doctoral degree in economics, and I would not have been able to write this dissertation without their input and advice. Sydney Ludvigson also deserves special thanks; it was a pleasure to work with her and Serena on the second chapter of this dissertation. During three of the years I spent working on this dissertation, I was funded by the National Science Foundation as a Graduate Research Fellow under grant number DGE-07-07425.

I am indebted to all the faculty members who contributed their comments and criticisms: Sally Davidson, Jennifer La'O, Emi Nakamura, Jaromir Nosal, José Scheinkman, Stephanie Schmitt-Grohé, Jón Steinsson, Martín Uribe, and Neng Wang; and to many of my classmates for their helpful discussions: Seungjun Baek, Ryan Chahrour, Nicolas Crouzet, Keshav Dogra, Ju Hyun Kim, Alexander McQuoid, Neil Mehrotra, Hyunseung Oh, Pablo Ottonello, Nikhil Patel, Dmitriy Sergeyev, Arunima Sinha, and Zigan Wang. I also thank my friends and family for their encouragement and prayers while I worked on this dissertation: my Redeemer community group; my good friends Toukam Ngoufanke, Jacob Tadros, and Jonathan Talamini; the Syvertsen family; the Paulson family; my abuelo Patricio, aunt Monica, and cousins Nicole, Nyah, and Naomi; my amazing parents Ed and Johanna; my beautiful sisters Kaila and Jenna; and my wonderful wife Ashley. Finally, I thank God for being my strength. From Him and through Him and for Him are all things. To Him be the glory forever.

To Abuela Alicia. I miss her.

Chapter 1

Advance Information and Distorted Beliefs in Macroeconomic and Financial Fluctuations

1.1 Introduction

An old idea in macroeconomics is that the beliefs of economic agents can be an important cause, and not merely the consequence, of developments in current economic conditions.¹ This causation can be proximate, as when agents receive advance information about future fundamentals. Or it can be ultimate, as when informational or psychological factors lead agents to hold beliefs that are distorted relative to the beliefs that would have been justified by fundamentals alone. Each of these possibilities has received renewed interest in recent years as a potentially important source of fluctuations in macroeconomic or financial data. However, most existing studies have not considered both possibilities at the same time. It therefore remains unclear whether one is more important than the other from a quantitative perspective. Or, given how similar the two hypotheses are, whether it is even possible to tell them apart from aggregate data alone.

This chapter argues that it is possible to separately identify the effects of advance infor-

¹The standard reference is [Pigou \(1927\)](#).

mation and distorted beliefs, and presents estimates of their relative importance in explaining fluctuations in macroeconomic and financial data. The estimates come from a medium-scale dynamic stochastic general equilibrium (DSGE) model populated by forward-looking agents whose beliefs may reflect advance information about future fundamentals or may be subject to exogenous distortions. Taken together, these two sources of exogenous variation in beliefs are responsible for two-thirds of the variation in the data. They are equally important for explaining output growth and the growth rate of hours worked. However, distorted beliefs are most important for explaining stock returns and consumption growth, and advance information is most important for explaining inflation and investment growth.

It is possible to separately identify advance information from distorted beliefs given data on endogenous forward-looking variables and the fundamentals relevant for determining those variables. At first, this may seem surprising, because agents can behave in exactly the same way whether they have received accurate good news about the future or whether they happen to be unjustifiably optimistic. What matters for their decisions today is *what* they believe about the future, not *why* they believe it. While that is true, the key for identification is that these two hypotheses generate different types of co-movement between current actions and the subsequent realizations of future fundamentals. The variation in current actions that results from advance information should be systematically related to future fundamentals, while the variation that results from distorted beliefs should not. With the benefit of hindsight, an econometrician can therefore examine the relationship between current actions and future fundamentals and determine the extent to which current actions are affected by these two different underlying causes.

In principle, any endogenous forward-looking variable should reveal information about the importance of advance information and distorted beliefs. Nevertheless, it is desirable to discipline the estimation procedure by incorporating data that reflect these types of belief-driven fluctuations as directly as possible. The two types of data selected in this study to satisfy that desideratum are subjective forecasts and stock returns. For the former,

agents' explicit reported forecasts of future activity are valuable because they represent a direct measure of agents' beliefs. For the latter, there is a large literature in financial economics suggesting that fluctuations in the stock market reflect agents' expectations about future activity. Indeed, the recent resurgence of advance information as a live hypothesis in macroeconomics has been based on that very idea.

A key methodological contribution of the chapter is in the way that belief distortions are specified. I propose to directly parameterize the Radon-Nikodým derivative process that governs the change of measure from agents' beliefs to the beliefs implied by their model environment. This approach has the theoretical advantage that it does not require a complete specification of the information that agents have at their disposal, or of the particular inferential procedure they use to update their beliefs. At a practical level, it circumvents the need to solve agents' (often complicated) filtering problems in order to obtain their equilibrium policy rules. It also makes it computationally straightforward to include belief distortions in a wide class of nonlinear equilibrium models. The idea of parameterizing the Radon-Nikodým process has been fruitfully employed in empirical asset pricing for handling stochastic risk adjustments, but has not yet been introduced into macroeconomic equilibrium models as a way of quantifying belief distortions.

The theoretical model developed in this chapter integrates several features that have been separately identified as important for explaining macroeconomic and financial data. These include six real frictions: recursive utility preferences, internal habit formation in consumption, investment adjustment costs, variable capacity utilization, and imperfect competition in labor and product markets. They also include two nominal frictions: time-dependent staggered price and wage setting. The real frictions are important for jointly explaining the behavior of stock returns and macroeconomic aggregates. The nominal frictions are important for explaining nominal wages, interest rates, and inflation. Together, these features help the model generate empirical predictions that are as consistent with the data as possible.

Solving and estimating models with recursive preferences introduces some technical chal-

lenges. To address those challenges this chapter presents a new framework for obtaining a linear approximation to the model that preserves the theoretically appealing features of this preference structure while maintaining the tractability that comes from working with a linear solution. The main insight underlying the method is that under the assumption of normal disturbances, the conditional expectation of a log-linear expression automatically incorporates information on conditional variances. By exploiting this well-known property of log-normal processes, it is possible to incorporate what is typically regarded as “second-moment” information into a linear approximation. The method is internally consistent in that it does not require approximating some equilibrium conditions differently than others. This is because the model itself naturally prioritizes the second-moment information that is particularly important for describing the dynamics in stochastic equilibrium. Furthermore, the method can be applied to many modern equilibrium models, including those with time-varying uncertainty.

This chapter makes a contribution to the literature on the sources of fluctuations in key macroeconomic and financial aggregates. If beliefs play an important role in the determination of those aggregates, it is of interest to understand what causes beliefs to change. In addition, there are normative reasons why the distinction between advance information and distorted beliefs can be important. For example, a monetary authority would want to respond differently to exogenous changes in private sector beliefs depending on which of these two sources were ultimately responsible. If it could tell the difference, it would want to accommodate advance information more than distorted beliefs, because the latter do not reflect any actual change in future economic fundamentals.² Of course telling the difference in real time can be difficult. During the stock market boom of the 1990s, former U.S. Federal Reserve Chairman Alan Greenspan highlighted this point in his now famous remark concerning the difficulty of determining whether “irrational exuberance” had unduly escalated asset values. While the positive analysis of this chapter does not attempt to solve this difficulty,

²Appendix A.4 provides a New-Keynesian example where this is the case. It is recommended to read section 1.3 before consulting this appendix.

it does bring us one step closer by providing a systematic way to tell the difference between these underlying causes after the fact.

As an overview of the chapter, section 1.2 discusses some studies that are related to this one. Section 1.3 describes the precise way that advance information and distorted beliefs are modeled, together with a simple permanent-income example to illustrate the main intuition behind identification. 1.4 presents the theoretical environment that serves as the laboratory for this chapter. Section 1.5 presents the solution and estimation methods that are applied to the model. Section 1.6 presents the data, estimates, and discusses the main results. Section 1.7 performs several robustness checks of these results. Section 1.8 explores how the model accounts for the joint behavior of the real economy and the stock market during two historical episodes: the Dot-Com boom and bust of the 1990s and the recent Great Recession. Section 1.9 and discusses useful extensions of the analysis in this chapter.

1.2 Related Literature

The idea that advance information about future fundamentals can be important for driving fluctuations in macroeconomic and financial data has received considerable attention in recent years. In the macroeconomic literature, Beaudry and Portier (2006) and Beaudry and Lucke (2010) use evidence from vector autoregressions (VARs) to argue that innovations to productivity growth are mostly anticipated, and have been a significant driving force of U.S. business cycles.³ Schmitt-Grohé and Uribe (2012) estimate a fully-specified DSGE model and find that advance information can account for half of the variation in output, consumption, investment, and employment. In the financial economic literature, Bansal and Yaron (2004) have pioneered a line of work that emphasizes the importance of “long-run risks” — advance information about future consumption growth — for explaining the notoriously

³I refrain from using the word “news” in this chapter to avoid any equivocation. In some contexts, “news” also refers to the realization of the *current* innovation to fundamentals, which would have been unforecastable one period before. I also refrain from using “long-run risks” because that label is less useful for emphasizing the informational content of this type of disturbance.

difficult risk-free rate, equity premium, and excess volatility puzzles.⁴

At the same time, exogenous distortions in agents’ beliefs have also played a prominent role in explaining business cycles and asset prices. Both [Angeletos and La’O \(2010\)](#) and [Lorenzoni \(2009\)](#) demonstrate how informational frictions can open the door for “noise shocks” — expectational errors arising from imperfect information about underlying fundamentals — to be an important cause of typical business cycle co-movement. In current ongoing work, [Angeletos et al. \(2014\)](#) present a model with dispersed information and heterogeneous priors where agents’ higher-order beliefs are subject to exogenous distortions. They find that this type of higher-order uncertainty, which they interpret as “confidence,” has the potential to explain a large fraction of business cycle fluctuations.⁵

Belief distortions have also played a prominent role in the asset-pricing literature. [Cecchetti et al. \(2000\)](#) argue that an exogenously distorted pessimism during expansions and optimism during contractions enables their consumption-based asset-pricing model to explain the first and second moments of the equity premium and risk-free rate puzzles. Connecting distorted beliefs about asset valuations to psychological factors has been the cornerstone of a large body of behavioral literature (e.g. [Shiller, 2015](#); [Akerlof and Shiller, 2009](#)). More recently, [Hassan and Mertens \(2014\)](#) emphasize the potential importance of belief distortions for the interaction between financial markets and the real economy. Their focus is on how small, correlated belief distortions about fundamentals can reduce the ability of stock prices to convey accurate information about the state of the economy.

A few studies explicitly consider advance information and distorted beliefs together as this chapter does. [Blanchard et al. \(2013\)](#) employ a similar information structure to [Lorenzoni \(2009\)](#) to identify the role of news and noise in macroeconomic data. While they find that noise appears to be a more important driver of business cycles, technically what is called

⁴For a more comprehensive discussion of the business cycle literature, see [Beaudry and Portier \(2014\)](#); for the long-run risks literature see [Bansal et al. \(2012\)](#).

⁵Appendix [A.2](#) discusses the type of disturbance proposed by [Angeletos et al. \(2014\)](#) along with two others that are closely related to the type of belief distortions described in the present chapter. Through simple examples, it clarifies the main conceptual differences and explains how those differences might be identified in the data. Incorporating direct data on subjective beliefs turns out to be important.

“news” in those papers does not represent what is called advance information here. This is because their news shock is not independent of current productivity. Instead, the key distinction in those papers is between permanent and transitory components of productivity. This chapter will feature advance information and distorted beliefs about both the permanent and transitory components of productivity, in addition to a number of other exogenous fundamental processes.

[Barsky and Sims \(2012\)](#) also make a distinction between advance information and distorted beliefs; they call the former “news” and the latter “animal spirits.” They use forward-looking measures of consumer confidence from the Michigan Survey of Consumers to help separately identify the two, and find that advance information is more important. By comparison, this chapter uses direct data on expectations from the Federal Reserve Bank of Philadelphia’s Survey of Professional Forecasters (SPF). Measures of consumer confidence may have the advantage of being representative of consumers’ beliefs more generally, with the disadvantage of being a less precise measure of expectations. Data from the SPF has also been more widely used for the purposes of estimating DSGE models.⁶

Another important difference with existing studies is that this chapter features an alternative formulation of distorted beliefs which does not require any explicit behavioral assumptions about how agents update their beliefs in the presence of new information.⁷ Specifically, I follow [Hansen and Sargent \(2005, 2007\)](#) and [Woodford \(2010\)](#) and characterize belief distortions according to the similarity of the probabilities that agents assign to future outcomes relative to the probabilities implied by their model environment. However, instead of using measures of statistical discrepancy to bound the degree of distortion, I propose to directly parameterize the Radon-Nikodým derivative process which summarizes the change from the distorted to historical probability measures. Given that the focus of this chapter is primarily positive rather than normative, this approach proves to be particularly useful both theo-

⁶For example, see [Del Negro and Eusepi \(2011\)](#), [Milani \(2011\)](#), and [Miyamoto and Nguyen \(2014\)](#).

⁷The models in the previous paragraph all assume that agents update their beliefs in a Bayesian manner. For models where agents use non-Bayesian procedures, see [Eusepi and Preston \(2011\)](#) and [Milani \(2011\)](#).

retically and computationally. The insight behind this added flexibility is that the goal of many learning models is essentially to solve for a particular change of measure, given an exogenously specified information structure. By specifying the change of measure directly, it is possible to avoid solving an often complicated intertemporal inference problem.

The idea of directly parameterizing belief distortions draws upon the practice in empirical asset pricing of parameterizing the change between the historical and *risk-neutral* measures. For example, this is the approach used in the term-structure models of [Dai and Singleton \(2000\)](#) and [Le et al. \(2010\)](#). This literature has shown that direct parameterization can be a useful way of uncovering many important properties of the equilibrium pricing kernel. In their study of bond risk premia, [Piazzesi et al. \(2013\)](#) also extend this approach to belief distortions as in this chapter. However, they do not focus on the distinction between advance information and distorted beliefs, and do not allow for independent exogenous variation in beliefs to arise through the Radon-Nikodým derivative process. They also do not apply this approach in a dynamic equilibrium setting, where optimizing agents choose their policy rules using a potentially distorted probability measure.

A related situation involving distorted beliefs is when agents fear that their model of the economy is misspecified. [Hansen et al. \(1999\)](#), [Cagetti et al. \(2002\)](#), and [Bidder and Smith \(2012\)](#) study equilibrium models where agents have a concern for robustness. Similarly, [Ilut and Schneider \(2014\)](#) endow agents with multiple priors utility to determine whether ambiguity about future fundamentals is responsible for business cycles. While these studies all involve belief distortions, those distortions reflect the degree to which agents are confident in *their own model* of the economy. According to the interpretation in this chapter, however, belief distortions do not necessarily arise due to any fear of model misspecification. Moreover, because distortions arising out of a concern for robustness reflect agents' beliefs about "worst case" outcomes, they may not be present in reported survey forecasts. Using forecast data therefore provides an important way of isolating distortions that are not the result of these

types of risk adjustments.⁸

In terms of the quantitative model, this chapter is related to a rapidly growing literature on the integration between asset markets and the broader macroeconomy in a general equilibrium setting. [Jermann \(1998\)](#) and [Boldrin et al. \(2001\)](#) both emphasize the importance of habit formation and real frictions in investment for jointly explaining asset returns and macroeconomic quantities in a production economy. More recently, a number of papers have focused on the role that recursive preferences can play in production-based economies, as opposed to habit formation; for example, [Kaltenbrunner and Lochstoer \(2010\)](#), [Croce \(2014\)](#), [Rudebusch and Swanson \(2012\)](#), [Gourio \(2012\)](#), and [Li and Palomino \(2014\)](#). One contribution of this chapter is to combine several different features from these studies. For example, it allows for both habits and recursive preferences, incorporates real frictions in production such as investment adjustment costs, and incorporates nominal rigidities and therefore a role for monetary policy. A second contribution is that the model in this chapter is estimated using likelihood-based methods. Most of the existing literature in this area has focused exclusively on calibration.

The solution method proposed in this chapter is closely related to the recent contribution of [Malkhozov \(2014\)](#). He also exploits the properties of log-normal random vectors to adjust policy functions for risk. However, when applied to homoskedastic models, his procedure only delivers a constant risk adjustment. That is, risk premia do not affect the impulse responses of endogenous variables to any of the exogenous disturbances. By contrast, the impulse response functions in this chapter depend on risk (and therefore risk preferences) because of a steady-state consistency criterion. Specifically, the steady state around which the approximation is constructed is required to coincide with the steady state implied by that approximation. Another difference is that the equilibrium conditions in this chapter are not separated into two groups before the solution procedure is applied — all equations are treated in the same way. Lastly, the presentation in this chapter explicitly derives the coefficients

⁸See [Appendix A.2](#) for further discussion of this point.

of the approximate solution from an arbitrary nonlinear model using matrix calculus. This facilitates comparison with other work such as [Schmitt-Grohé and Uribe \(2004\)](#).

1.3 Sources of Fluctuations in Beliefs

This section presents the mathematical details for how advance information and distorted beliefs are modeled. To highlight the intuition, the exposition is carried out through a familiar permanent income model of consumption and savings, which has only one fundamental process: income. In addition to being familiar, it has the advantage of admitting a closed-form solution, so that issues related to approximation can be postponed until Section 5. Nevertheless, it is important to keep in mind that this discussion applies to any forward-looking relation of the type described here; nothing is special about this consumption-based example in that respect. This section is meant to show three things: first, how both advance information and distorted beliefs can be included in an optimizing equilibrium model. Second, how both sources of fluctuations in beliefs can elicit similar responses from endogenous variables. And third, how it is nevertheless possible to separately identify the importance of these sources.

1.3.1 A Permanent-Income Example

The main distinctive aspect of this example (and the model of Section 1.4) is that it features two probability measures instead of one. To be precise about this distinction, let (Ω, \mathfrak{F}) denote the measurable space upon which the model economy is built. The set Ω contains all possible outcomes and \mathfrak{F} is a sigma-algebra of its subsets. On this space introduce two probability measures, both of which assign probabilities to events $\mathcal{A} \in \mathfrak{F}$. The first is the objective measure, denoted by \mathbb{P} , which defines the actual probability that any event will occur. The other is the distorted measure, denoted by \mathbb{Q} , which represents the subjective beliefs of economic agents within the model and may differ from the objective measure.

In the case that these two measures coincide, agents are said to have model-consistent (or “rational”) expectations.

The history of events in the economy is recorded by a sequence of non-decreasing sigma-algebras $\{\mathfrak{F}_t\}$ for $t \geq 0$. At each date, a representative agent chooses consumption c_t and debt holdings d_t so as to maximize his lifetime utility, which can be expressed recursively as:

$$V_t = u(c_t) + \beta E_t^{\mathbb{Q}}[V_{t+1}].$$

The expectation in this expression is the one induced by \mathbb{Q} , conditional on \mathfrak{F}_t . Note that the fact that \mathbb{Q} is a probability measure means that while the agent’s beliefs may be distorted, he nevertheless always maintains an internally-consistent set of beliefs. In particular, his subjective beliefs satisfy the law of iterated expectations and standard results in dynamic programming can be used to characterize his optimal actions.⁹

Each period, the agent receives a stochastic endowment of income, z_t . Exogenous innovations to z_t represent the “fundamental” disturbances driving the dynamics of the model. The agent is permitted to borrow or lend using a risk-free bond that pays a constant gross real interest rate $R > 1$. Under the simplifying assumptions of Hall (1978), namely that $u(c) = -\frac{1}{2}(c - \bar{c})^2$ for a bliss level of consumption $\bar{c} \geq 0$ and $\beta R = 1$, the optimal consumption choice satisfies¹⁰

$$c_t = (1 - \beta) \left[\sum_{\tau=0}^{\infty} \beta^{\tau} E_t^{\mathbb{Q}}[z_{t+\tau}] - R d_{t-1} \right]. \quad (1.1)$$

At each date, consumption equals the annuity value of expected total wealth: the present discounted value of expected lifetime income minus interest payments on outstanding debt. In order to complete the model, it is necessary to specify the evolution of z_t under \mathbb{P} and to clarify the relationship between \mathbb{P} and \mathbb{Q} .

⁹This differs, for example, from situations in which agents update their beliefs using simple econometric models. See Preston (2005) for a careful discussion of this point.

¹⁰As usual, this condition was derived under the assumption that when the agent decides on his debt holdings, he cannot expect to run a Ponzi scheme: $\lim_{j \rightarrow \infty} \beta^j E_t^{\mathbb{Q}}[d_{t+j}] \leq 0$.

1.3.2 Current Fundamentals

Without advance information or distorted beliefs, the only thing that can cause a change in beliefs is a change in current fundamentals. To see why this is the case, suppose that income is a first-order Markov process with independent, identically distributed normal innovations under the objective measure:

$$z_t = \rho_z z_{t-1} + e_t^z, \quad e_t^z \stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \sigma_z^2). \quad (1.2)$$

Under the hypothesis that the agent's beliefs are model-consistent (so that \mathbb{P} and \mathbb{Q} coincide), he forecasts future outcomes using the conditional distribution:

$$z_{t+1} | \mathfrak{F}_t \stackrel{\mathbb{P}}{\sim} \mathcal{N}(\rho_z z_t, \sigma_z^2).$$

The only source of variation in the right-hand side of this expression at date t relative to date $t-1$ comes from the contemporaneous innovation e_t^z , which enters through z_t . Because the income process is persistent, high current fundamentals signal high future fundamentals.

Under these assumptions, it is possible to solve (1.1) and show that consumption follows:

$$c_t = \left(\frac{1 - \beta}{1 - \beta \rho_z} \right) z_t - \left(\frac{1 - \beta}{\beta} \right) d_{t-1}.$$

In response to a positive innovation in income, part of the increase is consumed, part is used to pay off existing debt, and the rest is saved. Importantly, it follows that innovations in consumption only come as a result of innovations in current income. By introducing advance information and distorted beliefs, it is possible to break this tight link and allow changes in consumption to occur without any corresponding change in current income.

1.3.3 Advance Information

To allow for the possibility that the agent may be able to partially predict changes in future fundamentals, I introduce the new stochastic process a_t , which is related to z_t in the following way:¹¹

$$\begin{aligned} z_t &= \rho_z z_{t-1} + a_{t-1} + e_t^z \\ a_t &= \rho_a a_{t-1} + e_t^a, \quad e_t^a \stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \sigma_a^2). \end{aligned}$$

The disturbance e_t^a , which is independent and identically distributed over time and independent of e_t^z , represents advance information. It is unrelated to current fundamentals: $E^\mathbb{P}[e_t^a z_t] = 0$, but it is related to future fundamentals: $E^\mathbb{P}[e_t^a z_{t+\tau}] \neq 0$ for all $\tau > t$. The new conditional distribution relevant for forecasting future outcomes is:

$$z_{t+1} | \mathfrak{F}_t \stackrel{\mathbb{P}}{\sim} \mathcal{N}(\rho_z z_t + a_t, \sigma_z^2).$$

Now beliefs about z_{t+1} may fluctuate due to e_t^a as well as e_t^z , where the former disturbance enters through a_t . If $e_t^a = 0$ at all times, then no advance information about z_t is ever available and beliefs only fluctuate as a result of current developments.

With advance information, the equilibrium consumption rule from the permanent income model in (1.1) becomes:

$$c_t = \left(\frac{1 - \beta}{1 - \beta \rho_z} \right) \left[z_t + \left(\frac{\beta}{1 - \beta \rho_a} \right) a_t \right] - \left(\frac{1 - \beta}{\beta} \right) d_{t-1}.$$

Now, consumption may increase at date t either because the agent is wealthier due to an innovation in z_t , or because he has received advance information that he will be wealthier in

¹¹This dynamic structure is also present in Barsky and Sims (2011, 2012), Jinnai (2013), Barsky et al. (2014), Bansal and Yaron (2004), and Bansal et al. (2007, 2012). Alternative specifications include those by Christiano et al. (2010) and Schmitt-Grohé and Uribe (2012), who fix $\rho_a = 0$ and allow the process a_t to enter with a lag greater than one period: $z_t = \rho_z z_{t-1} + a_{t-\ell} + e_t^z$, with $\ell \geq 1$.

the future due to an innovation in a_t . The interesting aspect of modeling advance information in this way is that it allows the link between current actions and current fundamentals to be broken while still maintaining the model-consistency of beliefs. However, it still implies that beliefs are entirely determined by developments in fundamentals, once future developments are taken into account. Beliefs exert no genuinely ultimate causal influence on economic outcomes. The introduction of exogenous distortions in beliefs relaxes this restriction.

1.3.4 Distorted Beliefs

To introduce distortions in the agent's beliefs, I relax the restriction that subjective and objective probabilities must coincide. Specifically, I build \mathbb{Q} using a strictly positive \mathfrak{F}_t -measurable stochastic process M_t . At any date $t \geq 0$, \mathbb{Q} assigns probabilities to events $\mathcal{A} \in \mathfrak{F}_t$ according to the definition:

$$\mathbb{Q}(\mathcal{A}) \equiv E^{\mathbb{P}}[1_{\mathcal{A}}M_t],$$

where $1_{\mathcal{A}}$ is the indicator function for the set \mathcal{A} . In order for \mathbb{Q} to be a well-defined probability measure on the space (Ω, \mathfrak{F}) according to this definition, the process M_t must satisfy the following two criteria:

$$E_t^{\mathbb{P}}[M_{t+1}] = M_t \quad \text{and} \quad E^{\mathbb{P}}[M_t] = 1. \quad (1.3)$$

The first says that M_t is a martingale with respect to \mathbb{P} and the filtration $\{\mathfrak{F}_t\}_{t \geq 0}$. This is necessary so that \mathbb{Q} satisfies the Law of Iterated Expectations. The second guarantees that $\mathbb{Q}(\Omega) = 1$. The economic content of this specification of distortions is that the agent agrees with the model about which events are *impossible* (occur with probability zero), but he can arbitrarily disagree about how likely it is that any *possible* event may occur. The process M_t is known as a Radon-Nikodým derivative process, and represents the degree of disagreement between the agent and the model concerning these probabilities. In the case that $M_t = 1$ at

all times, there are no distortions, and the agent's beliefs are model-consistent.¹²

Even given these restrictions, the class of admissible distortions M_t is still too broad to be useful from the perspective of constructing economic models. One approach to restrict them further has been non-parametric: the extent to which subjective beliefs may be distorted relative to objective beliefs can be bounded on the basis of statistical discrepancy measures.¹³ An alternative approach, which is the one I pursue here, is to impose additional parametric structure on M_t . The parametric approach has been fruitfully applied in asset-pricing contexts to model risk-adjustments generated by stochastic discounting.¹⁴ For the example in this section, I assume that $M_0 = 1$ and that

$$M_{t+1} = M_t \exp \left(-\frac{1}{2} \frac{b_t^2}{\sigma_z^2} + \frac{b_t}{\sigma_z^2} e_{t+1}^z \right) \quad (1.4)$$

for an \mathfrak{F}_t -measurable stochastic process b_t . Notice that the first term inside the exponential guarantees that M_t is a martingale. Together with the initial condition, this implies that the required conditions in (1.3) are satisfied.

This parametric form is common in asset-pricing contexts and has a clear economic interpretation. Under the distorted measure induced by this process, it can be shown that the fundamental in (1.2) has the following dynamics under the distorted probability measure:

$$z_{t+1} | \mathfrak{F}_t \overset{\mathbb{Q}}{\sim} \mathcal{N}(\rho_z z_t + b_t, \sigma_z^2).$$

¹²This approach to modeling belief distortions is used, for example, by [Hansen and Sargent \(2005\)](#) in their work on robust control, and by [Woodford \(2010\)](#) in his formulation of “near-rational” expectations.

¹³See [Borovička et al. \(2014\)](#) for a discussion of the class of discrepancy measures that take the form:

$$\frac{1}{\theta(1+\theta)} E_t^{\mathbb{P}} \left[\left(\frac{M_{t+1}}{M_t} \right)^{1+\theta} - 1 \right],$$

and their usefulness for measuring the martingale component of stochastic discount factors. [Woodford \(2010\)](#) focuses on the limiting case of $\theta \rightarrow 0$, where this measure reduces to conditional relative entropy: $E_t^{\mathbb{P}}[M_{t+1}/M_t \ln(M_{t+1}/M_t)]$.

¹⁴See [Dai and Singleton \(2000\)](#) and [Le et al. \(2010\)](#) for discussions of affine term-structure models which rely on parameterizations of the stochastic discount factor.

That is, b_t represents a distortion in the conditional mean of the fundamental process. Higher values of b_t indicate that the agent is overly optimistic about future economic prospects relative to the model. Notice that under the restriction in (1.4), the agent still agrees with the model about the conditional variance of future random disturbances. A more general parametric class that also allows for distortions in the conditional variance is discussed in Section 1.5. However, there are two advantages to focusing on the conditional mean specification. First, because advance information is assumed only to affect conditional means and not conditional variances, this approach has the advantage of treating both sources of fluctuations in beliefs symmetrically; belief distortions resemble “mistaken” advance information. Second, with this restriction the model remains conditionally homoskedastic under both \mathbb{P} and \mathbb{Q} , which allows me to avoid some technical issues involved in solving models with time-varying volatility that are not the focus of this chapter.

Finally, it remains to specify a law of motion for the process b_t . Because it is the analogue of a_t , I endow it with the same dynamic structure:

$$b_t = \rho_b b_{t-1} + e_t^b, \quad e_t^b \stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \sigma_b^2).$$

The disturbance e_t^b is independent and identically distributed over time, independent of e_t^z , and captures genuinely exogenous variation in beliefs. Like advance information, it is unrelated to current fundamentals: $E^{\mathbb{P}}[e_t^b z_t] = 0$. But in contrast, it is unrelated to future fundamentals: $E^{\mathbb{P}}[e_t^b z_{t+\tau}] = 0$ for all $\tau > t$. With belief distortions of this type, the equilibrium consumption rule in (1.1) becomes

$$c_t = \left(\frac{1 - \beta}{1 - \beta \rho_z} \right) \left[z_t + \left(\frac{\beta}{1 - \beta \rho_b} \right) b_t \right] - \left(\frac{1 - \beta}{\beta} \right) d_{t-1}.$$

Consumption increases at date t either because the agent is wealthier today due to an innovation in z_t , or because he incorrectly believes that he will be wealthier in the future due to an innovation in b_t .

1.3.5 Combined System

Bringing together the three different potential sources of fluctuations in beliefs, I arrive at the following system of equations:

$$\begin{aligned} z_t &= \rho_z z_{t-1} + a_{t-1} + e_t^z, & e_t^z &\stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \sigma_z^2) \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(b_{t-1}, \sigma_z^2) \\ a_t &= \rho_a a_{t-1} + e_t^a, & e_t^a &\stackrel{\mathbb{P}, \mathbb{Q}}{\sim} \mathcal{N}(0, \sigma_a^2) \\ b_t &= \rho_b b_{t-1} + e_t^b, & e_t^b &\stackrel{\mathbb{P}, \mathbb{Q}}{\sim} \mathcal{N}(0, \sigma_b^2). \end{aligned} \tag{1.5}$$

Under these exogenous dynamics, an observed increase in consumption can come from one of three sources: an increase in current income, a correctly anticipated change in future fundamentals, or an incorrectly anticipated change in future fundamentals. For completeness, the policy rule combining (1.1) and (1.5) is:

$$c_t = \left(\frac{1 - \beta}{1 - \beta \rho_z} \right) \left[z_t + \left(\frac{\beta}{1 - \beta \rho_a} \right) a_t + \left(\frac{\beta}{1 - \beta \rho_b} \right) b_t \right] - \left(\frac{1 - \beta}{\beta} \right) d_{t-1}. \tag{1.6}$$

Before moving on, it is worth discussing some implicit restrictions imposed by this system, as well as the interpretation of the processes a_t and b_t . First, advance information and distorted beliefs are defined only relative to the fundamental process z_t . This means that there is no advance information about future belief distortions, and that there are no belief distortions about advance information. Moreover, there is also no advance information about future advance information, and there are no belief distortions about future belief distortions. While such multiple layers of fluctuations in beliefs would be theoretically straightforward to include, the more parsimonious structure employed here is sufficient for the purposes of this chapter. Nevertheless, that is technically a restriction that will be imposed throughout.

A second and related restriction is that the three innovations e_t^z , e_t^a , and e_t^b are independent. This rules out, for example, situations in which an innovation to current fundamentals *systematically* leads to a more than proportional increase in agents' optimism concerning fu-

ture fundamentals (whether that optimism is justified or not). In the context of a structural model such as the example in this section, it is sometimes possible to separately identify some forms of correlation between these disturbances. I will show how that can be done in the following subsection. The quantitative model in Section 1.4, however, will feature independent disturbances. This is in keeping with the business cycle literature, which typically assumes that all exogenous processes are independent.¹⁵

Third, it is important to understand the interpretation that is being given to the processes a_t and b_t . First consider b_t . This process is not meant to represent any variable that is “out there” in the model, which agents observe and incorrectly believe to be informative concerning future fundamentals. Rather, it is a useful mathematical way of representing the degree of correspondence between agents’ beliefs and the ones implied by the model, no matter how agents arrived at those beliefs.¹⁶ Similarly, on this interpretation the process a_t should not be interpreted as a variable that is literally observed by agents and known to be perfectly informative concerning future fundamentals. Instead, it is only meant to *reflect* that component of agents’ beliefs regarding the future which are systematically correlated with subsequent realizations of fundamental activity. Indeed, as I mentioned in the introduction, one advantage to modeling beliefs in this way is precisely to avoid making statements about what agents observe and what behavioral rules they use to update their beliefs based on that information.¹⁷

1.3.6 Identification: Key Moments

Given the similarity between advance information and distorted beliefs affect actions, how is it possible to separately identify their relative importance? To make this question precise,

¹⁵However, see Cúrdia and Reis (2010) for a discussion of why this might not be such a good idea, along with a methodological proposal for how to generalize that assumption in empirical work.

¹⁶See Woodford (2010) for a similar point. As in his analysis, the interpretation taken here implies that the usual assumption of model-consistent beliefs does not require that agents “know the model” and correctly solve its equations, but only that their beliefs coincide with the ones predicted by their model environment.

¹⁷Appendix A.3 discusses the type of belief distortions that can arise from making informational assumptions of these types.

consider an econometrician who understands the structure of the physical economy and is charged with estimating the relative importance of current fundamentals, advance information, and distorted beliefs. In terms of the the example of this section, he knows that in equilibrium consumption and debt holdings are related to income according to (1.6) and the budget constraint $c_t + d_{t-1}/\beta = z_t + d_t$, and that income and beliefs about income can be represented by the system (1.5). He also knows the value of β and the persistence parameters ρ_z , ρ_a , and ρ_b . Because he knows the persistence parameters, fix them at zero; this makes the key moment conditions especially stark. He is endowed with a sequence of observations of consumption and income generated by the model, and is charged with estimating the parameters σ_z , σ_a , and σ_b .

With a sample of observations of c_t and z_t , it is possible to separately identify these parameters. The intuition is that the three different disturbances induce different dynamic correlations between consumption and income. To isolate the key moments, note that in the case of no persistence, consumption growth satisfies:¹⁸

$$\Delta c_t = (1 - \beta) \left[z_t + \beta(a_t + b_t) - (a_{t-1} + b_{t-1}) \right].$$

First, the covariance between consumption growth and future income pins down the importance of advance information. Second, given σ_a , the variance of income pins down σ_z .¹⁹ Third, given σ_a and σ_z , the variance of consumption growth reveals the importance of distorted beliefs. Explicitly, these three moments are:

$$\begin{aligned} E^{\mathbb{P}}[\Delta c_t z_{t+1}] &= \beta(1 - \beta)\sigma_a^2 \\ E^{\mathbb{P}}[z_t^2] &= \sigma_a^2 + \sigma_z^2 \\ E^{\mathbb{P}}[\Delta c_t^2] &= (1 - \beta)^2 \left[\sigma_z^2 + \beta^2 \sigma_a^2 + (1 + \beta^2)\sigma_b^2 \right]. \end{aligned} \tag{1.7}$$

¹⁸Recall that consumption is not stationary in this model; it is a martingale under \mathbb{Q} .

¹⁹Note that it is also possible to use the covariance between consumption growth and income to identify σ_z (or to help identify β if it is also unknown): $E^{\mathbb{P}}[\Delta c_t z_t] = (1 - \beta)\sigma_z^2$.

In fact, it turns out that if the econometrician can directly observe data on subjective beliefs, it may be possible for him to identify these parameters even without any knowledge of the economic model. Suppose that his sample also includes observations of income forecasts, $E_t^{\mathbb{Q}}[z_{t+1}] = a_t + b_t$. Then he can infer the values of σ_z , σ_a , and σ_b from the three moments:

$$\begin{aligned} E^{\mathbb{P}}[E_t^{\mathbb{Q}}[z_{t+1}]z_{t+1}] &= \sigma_a^2 \\ E^{\mathbb{P}}[z_t^2] &= \sigma_a^2 + \sigma_z^2 \\ E^{\mathbb{P}}[E_t^{\mathbb{Q}}[z_{t+1}]^2] &= \sigma_a^2 + \sigma_b^2. \end{aligned} \tag{1.8}$$

Notice that this identification scheme does not require any knowledge of the structural relation between consumption and income (or any parameters in that relation, such as β). This highlights one of the main advantages of using direct evidence on subjective beliefs: in general, subjective forecasts of model variables provide additional moment restrictions without requiring any additional structural assumptions. This is because the dynamics of the structural model already contain predictions about the behavior of agents' forecasts. Later in this section, I will show in a Monte Carlo experiment that adding data on subjective forecasts can help to increase the precision of one's estimates of the relative importance of advance information and distorted beliefs.

In summary, the key moment for determining the relative importance of advance information is the covariance between a current endogenous variable (in this case either consumption or expected income) and future fundamentals. The key moment for determining the relative importance of current fundamentals is the variance of current fundamentals. And lastly, the key moment for determining the relative importance of distorted beliefs is the variance of the endogenous variable.

1.3.7 Correlations and Measurement Errors

This subsection briefly addresses two possible complications related to identifying the different sources of fluctuations in beliefs. First is the situation when either a_t or b_t is affected by current fundamental disturbances e_t^z ; in other words, the innovation to a_t or b_t is correlated with the current innovation to fundamentals. For example suppose that

$$b_t = \rho_b b_{t-1} + \tau e_t^z + e_t^b,$$

where $\tau \neq 0$. If $\tau > 0$ this implies that in response to a positive innovation in current fundamentals, agents' beliefs systematically become overly optimistic. In this case, is it possible to identify the parameter τ along with σ_z , σ_a , and σ_b ?

The answer is yes. The key additional moment is the correlation between the current endogenous variable and current fundamentals. In the consumption-income example of this section, still assuming for simplicity that $\rho_z = \rho_a = \rho_b = 0$ but now allowing for nonzero values of τ , it can be shown that

$$E^{\mathbb{P}}[\Delta c_t z_t] = (1 - \beta)(1 + \beta\tau)\sigma_z^2.$$

The mechanics behind the identification here is that σ_a and σ_z are already determined by the first two moments in system (1.7). If consumption growth is more correlated with income than these two parameters suggest based on the model, then it must be that the agent systematically overreacts to current fundamental developments. That is, that $\tau > 0$.

A second situation of interest is when variables are observed with error. To entertain this possibility, suppose that the econometrician only observes the following noisy measures

of consumption, income, and expected income:

$$\begin{aligned}\Delta\tilde{c}_t &= \Delta c_t + u_{\Delta c,t}, & u_{\Delta c,t} &\stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \omega_{\Delta c}^2) \\ \tilde{z}_t &= z_t + u_{z,t}, & u_{z,t} &\stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \omega_z^2) \\ \tilde{E}_t^{\mathbb{Q}}[z_{t+1}] &= E_t^{\mathbb{Q}}[z_{t+1}] + u_{E,t}, & u_{E,t} &\stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \omega_E^2).\end{aligned}$$

In this case, is it possible to identify the three parameters $\omega_{\Delta c}$, ω_z , and ω_E in addition to σ_z , σ_a , and σ_b ?

Again the answer is yes.²⁰ The reason is that classical measurement errors of the type assumed here only affect one observed series, and are uncorrelated with every other disturbance in the model. By contrast, each of the other three disturbances affect more than one variable. They generate particular patterns of *comovement* that are not affected by the measurement errors. Specifically, observe that the following six moments can be solved sequentially for all six unknown parameters:²¹

$$\begin{aligned}E^{\mathbb{P}}[E_t^{\mathbb{Q}}[z_{t+1}]z_{t+1}] &= \sigma_a^2 \\ E^{\mathbb{P}}[\Delta c_t z_t] &= (1 - \beta)\sigma_z^2 \\ E^{\mathbb{P}}[E_t^{\mathbb{Q}}[z_{t+1}]\Delta c_t] &= (1 - \beta)\beta(\sigma_a^2 + \sigma_b^2) \\ E^{\mathbb{P}}[z_t^2] &= \sigma_a^2 + \sigma_z^2 + \omega_z^2 \\ E^{\mathbb{P}}[E_t^{\mathbb{Q}}[z_{t+1}]^2] &= \sigma_a^2 + \sigma_b^2 + \omega_E^2 \\ E^{\mathbb{P}}[\Delta c_t^2] &= (1 - \beta)^2[\sigma_z^2 + \beta^2\sigma_a^2 + (1 + \beta^2)\sigma_b^2] + \omega_{\Delta c}^2\end{aligned}$$

This subsection has therefore demonstrated that it is in principle possible to separately identify the relative importance of distorted beliefs and advance information even in the

²⁰The answer is no in this example if the econometrician observes income and only one of either consumption or expected income. In general it is important that he observe more than one variable that is affected by belief distortions.

²¹In fact, one additional moment is still available either to identify β if it is unknown, or to allow for one additional correlation parameter. That moment is $E^{\mathbb{P}}[\Delta c_t z_{t+1}]$.

presence of some forms of correlation between the sources of fluctuations in beliefs, or classical measurement error. The next subsection presents an example of how these parameters can be estimated in practice.

1.3.8 Two Monte Carlo Experiments

The discussion of identification so far has focused on the case with no persistence in order to make the key moment conditions as stark as possible. In the more general case that the persistence parameters themselves must also be estimated from the data, each parameter can no longer be associated with just one moment condition. Instead, all the parameters must be estimated jointly from a set of moment conditions. The approach of this chapter is to prioritize different moments according to the way they are encoded in the likelihood function. This subsection presents two Monte Carlo experiments that illustrate how the different parameters of the consumption-income example can be identified based on the likelihood function. It also illustrates the additional precision afforded by exploiting subjective forecasts.

The true parameters are fixed at $(\rho_z, \rho_a, \rho_b, \sigma_z, \sigma_a, \sigma_b) = (0.9, 0.7, 0.5, 0.4, 0.6, 0.8)$. The subjective discount factor is assumed to be known and is set to $\beta = 0.995$. Given these true parameters, I simulate 10,000 artificial samples of consumption growth Δc_t , income growth Δz_t , and expected income growth $E_t^Q[\Delta z_{t+1}]$ from the consumption-income example in this section. Each sample consists of 219 observations of each variable, which is the same size as the actual data sample that will be used in Section 1.6. All parameter estimates are obtained by numerically maximizing the likelihood function implied by the model.

In the first experiment, suppose that the econometrician only observes consumption growth and income growth but no subjective forecasts. The dashed lines in Figure 1.1 represent the distribution of parameter estimates based on these two variables. In each case the mean is close to the true value (represented by the vertical dotted line), indicating that the likelihood function does a relatively good job at revealing the values of these parameters

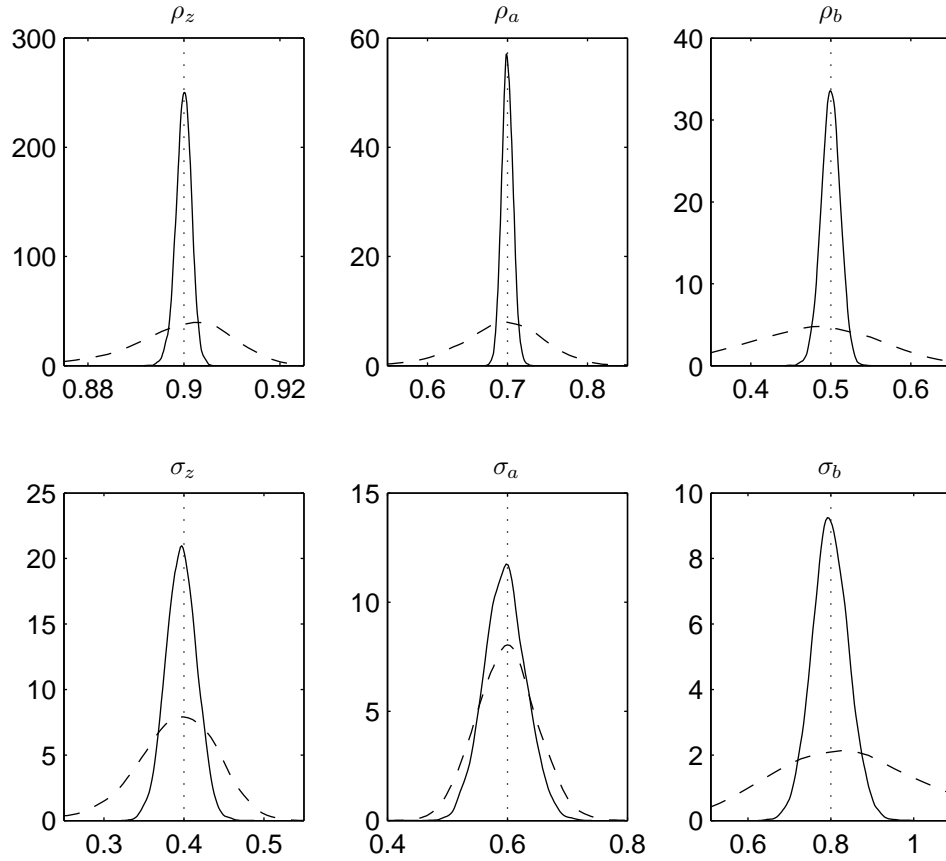


Figure 1.1: Monte Carlo experiments. This figure shows the distribution of parameter estimates from the consumption-income example discussed in Section 1.3. In each subplot, the vertical dotted line indicates the true value of the parameter. Based on these values, 10,000 samples of 219 observations are simulated for estimation. Estimates are obtained by numerically maximizing the likelihood function implied by the model. The dashed line is the distribution of estimates based only on observations of consumption growth and income growth. The solid line is the distribution when observations of subjective forecasts of income growth are included in addition to consumption growth and income growth. The first column of panels corresponds to the parameters related to contemporaneous innovations in income (persistence ρ_z and conditional volatility σ_z), the second column to the parameters related to advance information about future income (ρ_a, σ_a), and the third to the parameters related to distorted beliefs about future income (ρ_b, σ_b).

for a sample of this size. Note however that the parameters related to belief distortions (the third column) are the most imprecisely estimated.

In the second experiment, suppose that the econometrician is also given data on subjective forecasts. In particular, he observes subjective expectations of future income growth in addition to consumption growth and income growth. The solid lines in Figure 1.1 represent the distribution of parameter estimates based on these three variables. The main difference between this and the first experiment is clear: by including data on subjective forecasts, the precision of estimates notably increases. The distributions are still centered around their true values, but now the dispersion of the distribution is reduced, especially for the persistence parameters and for the conditional variance of belief distortions.

1.4 Quantitative Equilibrium Model

This section presents the central model of the chapter. It has three distinctive features. First, agents have beliefs that may reflect advance information or exogenous distortions about future fundamentals. Second, “fundamentals” are defined by a rich set of nine exogenous processes. This is to ensure that advance information and distorted beliefs are not merely standing in for other omitted processes commonly considered in the literature. Third, households have recursive utility preferences, which are included to help the model jointly explain macroeconomic quantities and asset returns.

1.4.1 Household

The model features a representative household with preferences that satisfy the recursion:²²

$$V_t = \left\{ (1 - \beta) \zeta_t^C U_t^{1-1/\xi} + \beta E_t^{\mathbb{Q}} [V_{t+1}^{1-\gamma}]^{\frac{1-1/\xi}{1-\gamma}} \right\}^{\frac{1}{1-1/\xi}}. \quad (1.9)$$

²²Preferences of this type are also often called “Epstein-Zin” or “Epstein-Zin-Weil” preferences in reference to [Epstein and Zin \(1989\)](#) and [Weil \(1989\)](#). The original framework for recursive utility preferences was developed by [Kreps and Porteus \(1978\)](#).

The variable $U_t \geq 0$ denotes the period utility kernel, and ζ_t^C is an exogenous stochastic disturbance to preferences. The parameter $\beta \in (0, 1)$ determines the marginal rate of time preference, $\gamma > 0$ controls the degree of risk aversion, and $\xi > 0$ controls the degree to which the household is willing to substitute period utility flows over time.²³ In the special case that $\gamma = 1/\xi$, these preferences collapse to the standard power utility form. If $\gamma > 1/\xi$, the household has a preference for early resolution of uncertainty, and if $\gamma < 1/\xi$ it has a preference for late resolution. As in the model of the previous section, expectations are computed according to the distorted probability measure \mathbb{Q} , which is common to all agents in the model.

The period utility kernel U_t is a function of current and past consumption, C_t , and hours worked H_t :

$$U_t = C_t - \phi C_{t-1} - \vartheta H_t^\psi D_t. \quad (1.10)$$

The variable D_t is a geometric average of current and past habit-adjusted consumption.²⁴ Its law of motion is assumed to be:

$$D_t = (C_t - \phi C_{t-1})^\nu D_{t-1}^{1-\nu}. \quad (1.11)$$

The parameter $\phi \in [0, 1)$ controls the degree of internal habit formation in preferences, $\psi > 1$ determines the Frisch elasticity of labor supply when $\nu = \phi = 0$, and $\vartheta > 0$ scales the disutility of labor.

These preferences have two appealing features. First, they nest as special cases two widely used utility functions in the business cycle literature. If $\nu = 1$, $\phi = 1$, and $\gamma = 1/\xi$, preferences are of the type considered by [King et al. \(1988\)](#); if $\nu = 0$, $\phi = 1$, and $\gamma = 1/\xi$,

²³When $\xi = 1$ preferences are defined by the pointwise limit:

$$V_t = \left\{ U_t^{(1-\beta)\zeta_t^C} E_t^{\mathbb{Q}} \left[V_{t+1}^{1-\gamma} \right]^{\frac{\beta}{1-\gamma}} \right\}^{\frac{1}{(1-\beta)\zeta_t^C + \beta}}.$$

²⁴The introduction of this process is due to [Jaimovich and Rebelo \(2009\)](#), and [Schmitt-Grohé and Uribe \(2012\)](#) modify it to incorporate habit persistence.

there are no wealth effects on labor supply as in [Greenwood et al. \(1988\)](#). Second, they nest as special cases two widely used utility functions in the asset pricing literature. If $\phi > 0$, $\gamma = 1/\xi$, and $\vartheta = 0$, utility is derived from current consumption relative to a (one period) habit stock of consumption and the household is indifferent concerning the intertemporal resolution of uncertainty as in [Campbell and Cochrane \(1999\)](#). If $\phi = 0$, $\gamma \neq 1/\xi$, and $\vartheta = 0$, there are no habits but the household is not indifferent about how uncertainty is resolved as in [Bansal and Yaron \(2004\)](#). A byproduct of the estimation will therefore be to determine the quantitative impact of each of these four prominent features of preferences.

The household has access to a complete set of Arrow securities, each of which represents a claim to a real state-contingent return in the subsequent period. The aggregate resources devoted to these securities at date t is denoted by A_t , and they deliver the real state-contingent return R_{t+1}^A at date $t + 1$. The household also receives wage income W_t^s from supplying its raw labor services to a monopolistic union. Letting T_t denote net lump-sum transfers, the household's flow budget constraint is:

$$C_t + A_t \leq W_t^s H_t + R_t^A A_{t-1} + T_t.$$

The absence of arbitrage opportunities in this economy requires that there exist a unique stochastic discount factor S_{t+1} with the property that the state-contingent return on any asset must satisfy the relation $E_t^Q[S_{t+1} R_{t+1}^A] = 1$. Letting $\Lambda_t U_t^{-1/\xi}$ denote the household's marginal utility of income, the real stochastic discount factor in this economy is:

$$S_{t+1} = \beta \frac{\Lambda_{t+1}}{\Lambda_t} \left(\frac{U_{t+1}}{U_t} \right)^{-1/\xi} \left(\frac{V_{t+1}}{E_t^Q[V_{t+1}^{1-\gamma}]^{1/(1-\gamma)}} \right)^{1/\xi-\gamma}. \quad (1.12)$$

Between t and $t + \tau$, the stochastic discount factor is $S_{t,t+\tau} \equiv \prod_{s=1}^{\tau} S_{t+s}$.

When there are no habits in consumption ($\phi = 0$), the utility kernel is separable in consumption and leisure ($\nu = 0$), and there is no exogenous variation in preferences ($\zeta_t^C \equiv 1$), the marginal utility of income is equal to $U_t^{-1/\xi}$. On the other hand when those features are

present, the marginal utility of income depends on how income will affect the habit stock and disutility of labor in subsequent periods. This forward-looking aspect can be seen in the optimality condition with respect to consumption:

$$\Lambda_t = (1 - \beta)\zeta_t^C - \nu\Gamma_t \frac{D_t}{C_t - \phi C_{t-1}} - \phi E_t^{\mathbb{Q}} \left[S_{t+1}^* \left((1 - \beta)\zeta_{t+1}^C - \nu\Gamma_{t+1} \frac{D_{t+1}}{C_{t+1} - \phi C_t} \right) \right], \quad (1.13)$$

where $S_{t+1}^* \equiv S_{t+1}\Lambda_t/\Lambda_{t+1}$. The household internalizes the effects of its consumption choices on the stock variable D_t . The marginal value of increasing this variable by one unit is denoted $\Gamma_t U_t^{-1/\xi}$, where Γ_t follows:

$$\Gamma_t = \vartheta\psi H_t^{\psi}(1 - \beta)\zeta_t^C + (1 - \nu)E_t^{\mathbb{Q}} \left[S_{t+1}^* \Gamma_{t+1} \frac{D_{t+1}}{C_{t+1} - \phi C_t} \right]. \quad (1.14)$$

Lastly, the household takes wages as given, and supplies labor up to the point where the ratio of its marginal disutility of work to its marginal utility of income equals the real wage:

$$\frac{\vartheta\psi H_t^{\psi-1} D_t (1 - \beta)\zeta_t^C}{\Lambda_t} = W_t^s. \quad (1.15)$$

1.4.2 Labor Union

A centralized labor union receives raw labor input H_t at a real marginal cost of $\zeta_t^W W_t^s$, and specializes it for use in a continuum of labor markets indexed by $j \in [0, 1]$. The union supplies this specialized labor monopolistically, and ζ_t^W captures exogenous variation in marginal costs that creates a wedge between the wage it charges in labor market j and the wage it pays to the household. In each labor market, the union supplies enough specialized labor to satisfy demand:

$$H_t(j) = \left(\frac{W_t(j)}{W_t} \right)^{-\eta_w} H_t^d.$$

The variable $W_t(j)$ denotes the real wage charged by the union in labor market j , W_t is an index of wages in the economy, and H_t^d represents the aggregate level of specialized labor demanded by productive firms.

In each period, the union is only capable of resetting wages optimally in a randomly chosen fraction $1 - \theta_W \in [0, 1)$ of labor markets.²⁵ In the remaining θ_W markets, nominal wage growth is indexed to past gross inflation according to rule:

$$\frac{W_t(j)}{W_{t-1}(j)} \Pi_t = (\mu_W \Pi_{t-1})^{\chi_W}.$$

The constant term μ_W is the steady-state growth rate of real wages, and is determined in equilibrium. The variable Π_t denotes the gross inflation rate from period $t - 1$ to t . The parameter $\chi_W \in [0, 1]$ controls the degree of wage indexation, with $\chi_W = 0$ indicating no indexation.

In each labor market where the union is able to optimally set its posted wage, it maximizes the present discounted value of future real profits,

$$E_t^Q \sum_{\tau=0}^{\infty} \theta_W^\tau S_{t,t+\tau} \left\{ (W_{t+\tau}(j) - \zeta_{t+\tau}^W W_{t+\tau}^s) H_{t+\tau}(j) \right\},$$

subject to the demand curve and indexation rule above. All profits are re-distributed lump sum to households, so state-contingent profit streams are valued using the stochastic discount factor $S_{t,t+\tau}$. Marginal costs are common across all labor markets, so the union will post the same wage W_t^* and supply the same amount of labor H_t^* in all markets where it is able to re-optimize. Optimality requires that W_t^* equates the expected marginal gain of having that wage persist forever in terms of additional revenue per unit of labor demand to a constant

²⁵For a discussion of the differences between this approach to modeling nominal wage rigidity and other common approaches such as the one in [Erceg et al. \(2000\)](#), see [Schmitt-Grohé and Uribe \(2006\)](#). The main advantage here is that the assumption of a centralized market for the raw household labor input ensures that equilibrium heterogeneity in labor supply does not also lead to heterogeneity in consumption.

markup over the marginal loss in terms of reduced demand:

$$\frac{W_t^* E_t^{\mathbb{Q}} \sum_{\tau=0}^{\infty} \theta_W^{\tau} S_{t,t+\tau} H_{t+\tau}^d W_{t+\tau}^{\eta_W} \left(\prod_{s=1}^{\tau} \frac{(\mu_W \Pi_{t+s-1})^{\chi_W}}{\Pi_{t+s}} \right)^{1-\eta_W}}{E_t^{\mathbb{Q}} \sum_{\tau=0}^{\infty} \theta_W^{\tau} S_{t,t+\tau} H_{t+\tau}^d W_{t+\tau}^{\eta_W} \zeta_{t+\tau}^W W_{t+\tau}^s \left(\prod_{s=1}^{\tau} \frac{(\mu_W \Pi_{t+s-1})^{\chi_W}}{\Pi_{t+s}} \right)^{-\eta_W}} = \left(\frac{\eta_W}{\eta_W - 1} \right).$$

In the case that $\theta_W \rightarrow 0$, this condition says that the optimal reset wage is equal to a time-varying markup over the wage paid to the household:

$$W_t^* = \left(\frac{\eta_W}{\eta_W - 1} \right) \zeta_t^W W_t^s.$$

It is possible to express this optimal pricing condition recursively by introducing an auxiliary forward-looking variable Υ_t^W . This variable satisfies the two relations:

$$\Upsilon_t^W = \left(\frac{\eta_W}{\eta_W - 1} \right) \zeta_t^W W_t^s H_t^d W_t^{\eta_W} + \theta_W E_t^{\mathbb{Q}} \left[S_{t+1} \left(\frac{(\mu_W \Pi_t)^{\chi_W}}{\Pi_{t+1}} \right)^{-\eta_W} \Upsilon_{t+1}^W \right] \quad (1.16)$$

$$\Upsilon_t^W = W_t^* H_t^d W_t^{\eta_W} + \theta_W E_t^{\mathbb{Q}} \left[S_{t+1} \left(\frac{(\mu_W \Pi_t)^{\chi_W}}{\Pi_{t+1}} \right)^{1-\eta_W} \frac{W_t^*}{W_{t+1}^*} \Upsilon_{t+1}^W \right] \quad (1.17)$$

1.4.3 Productive Firm

To produce wholesale goods, a representative firm hires labor, accumulates capital, and controls how intensely it uses its currently installed capital. Its production function is:

$$Y_t = \zeta_t^Y (u_t K_{t-1})^{\alpha_K} (Z_t H_t^d)^{\alpha_H} (Z_t L)^{1-\alpha_K-\alpha_H}. \quad (1.18)$$

The variable ζ_t^Y is an exogenous, transitory productivity process, and Z_t is a permanent one. A fixed amount of land L is used in production, which generates decreasing returns to scale in the other factors of production. In the special case that $\alpha_K + \alpha_H = 1$, this production function exhibits constant returns to scale in capital and labor. Variable capacity utilization is captured by u_t .

The labor input used by the firm is a composite of labor from all different labor markets:

$$H_t^d = \left[\int_0^1 H_t(j)^{1-1/\eta_W} dj \right]^{1/(1-1/\eta_W)},$$

where $\eta_W > 1$ controls the degree of substitutability across different labor types. The labor demand function faced by the centralized union is the solution to the firm's problem of minimizing total labor cost subject to this aggregation technology.

Capital accumulation is subject to convex investment adjustment costs:

$$K_t = (1 - \delta(u_t))K_{t-1} + \left[1 - \Phi\left(\frac{I_t}{I_{t-1}}\right) \right] \zeta_t^I I_t, \quad (1.19)$$

where ζ_t^I represents an exogenous disturbance in the rate of transformation from investment goods I_t to capital goods, and $\delta(u_t) \equiv \delta_0 + \delta_1(u_t - 1) + \delta_2/2(u_t - 1)^2$ is the stochastic rate of depreciation, which increases if the firm employs its existing capital stock more intensely. Adjustment costs are quadratic in deviations of investment growth from its steady-state level: $\Phi(I_t/I_{t-1}) \equiv \frac{\kappa}{2}(I_t/I_{t-1} - \mu_I)^2$ with $\kappa > 0$. The total cost of investment goods in terms of the consumption good is $X_t I_t$, where X_t is a non-stationary exogenous process that represents the “relative price of investment.”²⁶

Output is sold in a competitive, centralized market at price P_t^s . Profits earned by the firm are rebated lump sum to the household. Its objective is to maximize the present discounted value of its profits from production, where future payouts are valued according to the household's stochastic discount factor. The (cum-dividend) value of the firm obeys the recursion:

$$J_t = P_t^s Y_t - W_t H_t^d - X_t I_t + E_t^{\mathbb{Q}}[S_{t+1} J_{t+1}]. \quad (1.20)$$

Equity in this model is defined as a levered claim on the productive firm. Letting $\lambda \geq 1$

²⁶Alternatively, because X_t captures changes in the ability of firms to transform investment goods into consumption goods, it is sometimes also referred to as “investment-specific technological progress” or “capital-embodied technological progress.”

capture the degree of leverage, the gross real return on equity is therefore a function of the market return on the firm:²⁷

$$R_{t+1}^e = \left(\frac{J_{t+1}}{J_t} \right)^\lambda. \quad (1.21)$$

Turning to the optimal decisions of the firm, labor is demanded up to the point that its marginal product is equals the aggregate wage:

$$\alpha_H \frac{P_t^s Y_t}{H_t^d} = W_t, \quad (1.22)$$

and the desired level of capital must satisfy the no-arbitrage condition:

$$1 = E_t^{\mathbb{Q}} \left[S_{t+1} \left(\frac{\alpha_K Y_{t+1}/K_t + Q_{t+1}(1 - \delta(u_{t+1}))}{Q_t} \right) \right], \quad (1.23)$$

where the term in parentheses is the gross rate of return on investment. The variable Q_t represents (marginal) Tobin's Q — the value of installed capital in terms of its replacement cost. In the absence of any adjustment costs or any transitory investment-specific technological change ($\zeta_t^I \equiv 1$) Tobin's Q is equal to the replacement cost of capital in terms of the consumption good: $Q_t = X_t$. In the presence of adjustment costs however, the replacement cost of capital also depends on the growth rate of investment:

$$\begin{aligned} X_t = \zeta_t^I Q_t & \left[1 - \Phi \left(\frac{I_t}{I_{t-1}} \right) - \Phi' \left(\frac{I_t}{I_{t-1}} \right) \frac{I_t}{I_{t-1}} \right] \\ & + E_t^{\mathbb{Q}} \left[S_{t+1} \zeta_{t+1}^I Q_{t+1} \Phi' \left(\frac{I_{t+1}}{I_t} \right) \left(\frac{I_{t+1}}{I_t} \right)^2 \right]. \end{aligned} \quad (1.24)$$

Lastly, the level of intensity at which the firm utilizes its installed capital equates the marginal

²⁷Related studies such as [Kaltenbrunner and Lochstoer \(2010\)](#) and [Croce \(2014\)](#) instead associate equity returns with levered returns to a capital claim. With constant returns to scale, these two approaches are the same (for example, see [Restoy and Rockinger, 1994](#)). However, the model in this chapter allows for non-constant returns to scale; therefore this definition is the appropriate one.

benefit of higher output and the marginal cost of a higher rate of depreciation:

$$\alpha_K \frac{Y_t}{K_{t-1}} = Q_t u_t \delta'(u_t). \quad (1.25)$$

1.4.4 Retail Firm

A monopolistic, centralized retailer purchases wholesale goods from productive firms at a per-unit cost $\zeta_t^P P_t^s$, and specializes them for sale in a continuum of retail markets indexed by $i \in [0, 1]$. The variable ζ_t^P captures exogenous variation in marginal costs and creates a wedge between the price that the retailer charges in market i and the price it pays the productive firm. The demand curve in each market is given by

$$Y_t(i) = P_t(i)^{-\eta_P} Y_t^d,$$

where $P_t(i)$ denotes the price posted in market i relative to an aggregate price index, and Y_t^d is the aggregate amount of final goods demanded in the economy.

As with the labor union, the retailer is only capable of optimally resetting wages in a randomly chosen fraction $1 - \theta_P \in [0, 1]$ of product markets. In the remaining θ_P markets, nominal product price growth is indexed to past inflation according to the rule:

$$\frac{P_t(i)}{P_{t-1}(i)} \Pi_t = \Pi_t^{\chi_P},$$

where $\chi_P \in [0, 1]$ controls the degree of price indexation.

In each product market the retailer is able to optimally adjust its posted price, it maximizes the present discounted value of all future real profits,

$$E_t^Q \sum_{\tau=0}^{\infty} \theta_P^\tau S_{t,t+\tau} \left\{ (P_{t+\tau}(i) - \zeta_{t+\tau}^P P_{t+\tau}^s) Y_{t+\tau}(i) \right\},$$

subject to the demand curve and indexation rule above. All profits are redistributed lump

sum to households. Because marginal costs are common across all product markets, the price will be the same in all markets that are re-optimized in period t . This price, P_t^* , satisfies the optimality condition:

$$\frac{P_t^* E_t^{\mathbb{Q}} \sum_{\tau=0}^{\infty} \theta_P^{\tau} S_{t,t+\tau} Y_{t+\tau}^d \left(\prod_{s=1}^{\tau} \frac{(\Pi_{t+s-1})^{\chi_P}}{\Pi_{t+s}} \right)^{1-\eta_P}}{E_t^{\mathbb{Q}} \sum_{\tau=0}^{\infty} \theta_P^{\tau} S_{t,t+\tau} Y_{t+\tau}^d \zeta_{t+\tau}^P P_{t+\tau}^s \left(\prod_{s=1}^{\tau} \frac{(\mu_P \Pi_{t+s-1})^{\chi_P}}{\Pi_{t+s}} \right)^{-\eta_P}} = \left(\frac{\eta_P}{\eta_P - 1} \right).$$

This optimality condition can also be expressed recursively by introducing the auxiliary forward-looking variable Υ_t^P , which satisfies the two relations:

$$\Upsilon_t^P = \left(\frac{\eta_P}{\eta_P - 1} \right) \zeta_t^P P_t^s Y_t^d + \theta_P E_t^{\mathbb{Q}} \left[S_{t+1} \left(\frac{\Pi_t^{\chi_P}}{\Pi_{t+1}} \right)^{-\eta_P} \Upsilon_{t+1}^P \right] \quad (1.26)$$

$$\Upsilon_t^P = P_t^* Y_t^d + \theta_P E_t^{\mathbb{Q}} \left[S_{t+1} \left(\frac{\Pi_t^{\chi_P}}{\Pi_{t+1}} \right)^{1-\eta_P} \frac{P_t^*}{P_{t+1}^*} \Upsilon_{t+1}^P \right]. \quad (1.27)$$

1.4.5 Government

The government consumes an exogenous amount of goods, G_t , each period. The fiscal authority finances that consumption through lump-sum taxation of the household, and maintains a balanced budget each period. The monetary authority targets the return on nominally risk-free bonds according to a feedback rule of the form:

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R} \right)^{\varphi_R} \left[\left(\frac{\Pi_t}{\Pi} \right)^{\varphi_{\Pi}} \left(\frac{Y_t^d / Y_{t-1}^d}{\mu_Y} \right)^{\varphi_{Y_d}} \right]^{1-\varphi_R} \zeta_t^R. \quad (1.28)$$

The nominal rate is a function of deviations of gross inflation from its long-run target level Π , deviations of aggregate demand growth from its steady state level μ_Y , and potentially also past rates. The variable ζ_t^R captures exogenous variation in the nominal rate.

By no arbitrage, the equilibrium ex-post real interest rate between period t and $t+1$, denoted R_t / Π_{t+1} , must satisfy the pricing condition:

$$1 = E_t^{\mathbb{Q}} \left[S_{t+1} \frac{R_t}{\Pi_{t+1}} \right]. \quad (1.29)$$

1.4.6 Market Clearing and Equilibrium

There are four non-financial markets in this economy: the markets for unspecialized labor, specialized labor, wholesale goods, and retail goods. Market clearing requires that aggregate supply equal aggregate demand in each of these markets.

The amount of unspecialized labor supplied by households must equal the amount demanded by the union, $H_t = \int_0^1 H_t(j) dj$, and in turn the amount of specialized labor supplied by the union must equal the amount demanded by productive firms, $\int_0^1 H_t(j) dj = H_t^d \int_0^1 (W_t(j)/W_t)^{-\eta_w} dj$. Defining $\Delta_t^W \equiv \int_0^1 (W_t(j)/W_t)^{-\eta_w} dj$, these relations imply that a fraction of the labor supplied by households is lost due to the inefficient wage dispersion:

$$H_t = \Delta_t^W H_t^d. \quad (1.30)$$

Similarly, a fraction of the total output supplied by productive firms is lost as a result of inefficient price distortions. Defining terms analogously to the labor market:

$$Y_t = \Delta_t^P Y_t^d, \quad (1.31)$$

where the aggregate demand for final goods is given by

$$Y_t^d = C_t + X_t I_t + G_t. \quad (1.32)$$

The nature of the time-dependent rigidities in these markets allows the laws of motion for the aggregate wage index, price index, wage dispersion, and price dispersion to be written in recursive form. In that order, the laws of motion are:

$$W_t^{1-\eta_w} = \theta_W \left(\frac{(\mu_W \Pi_t)^{\chi_W}}{\Pi_{t+1}} \right)^{1-\eta_w} W_{t-1}^{1-\eta_w} + (1 - \theta_W) (W_t^*)^{1-\eta_w} \quad (1.33)$$

$$1 = \theta_P \left(\frac{\Pi_{t-1}^{\chi_P}}{\Pi_t} \right)^{1-\eta_P} + (1 - \theta_P) (P_t^*)^{1-\eta_P} \quad (1.34)$$

$$\Delta_t^W = \theta_W \left(\frac{W_{t-1} (\mu_W \Pi_{t-1})^{\chi_W}}{W_t \Pi_t} \right)^{-\eta_W} \Delta_{t-1}^W + (1 - \theta_W) \left(\frac{W_t^*}{W_t} \right)^{-\eta_W} \quad (1.35)$$

$$\Delta_t^P = \theta_P \left(\frac{\Pi_{t-1}^{\chi_P}}{\Pi_t} \right)^{-\eta_P} \Delta_{t-1}^P + (1 - \theta_P) (P_t^*)^{-\eta_P}. \quad (1.36)$$

Having laid out all the necessary conditions, I can now define the competitive equilibrium:

Definition 1. *A competitive equilibrium in the model from Section 1.4 is a set of stochastic processes $\{C_t, I_t, Y_t, Y_t^d, H_t, H_t^d, K_t, U_t, D_t, Q_t, V_t, J_t, \Lambda_t, \Gamma_t, u_t, S_t, R_t, R_t^e, W_t^s, W_t^*, W_t, P_t^s, P_t^*, \Pi_t, \Upsilon_t^W, \Upsilon_t^P, \Delta_t^P, \Delta_t^W\}_{t=0}^\infty$ satisfying the system of equations (1.9)-(1.36), given the set of exogenous stochastic processes $\{\zeta_t^C, \zeta_t^I, \zeta_t^Y, \zeta_t^W, \zeta_t^P, \zeta_t^R, G_t, Z_t, X_t\}_{t=0}^\infty$ and initial conditions.*

1.4.7 Fundamental Processes

The model features nine fundamental processes, representing exogenous variation in permanent neutral productivity (Z_t), transitory neutral productivity (ζ_t^Y), permanent investment-specific productivity (X_t), transitory investment-specific productivity (ζ_t^I), labor market marginal costs (ζ_t^W), product market marginal costs (ζ_t^P), preferences (ζ_t^C), government spending (G_t), and monetary policy (ζ_t^R). The two permanent shocks are assumed to have stationary growth rates:

$$\mu_{Z,t} \equiv \frac{Z_t}{Z_{t-1}} \quad \text{and} \quad \mu_{X,t} \equiv \frac{X_t}{X_{t-1}}.$$

These two permanent shocks generate trend growth in several endogenous variables. The trends in output and investment, for example are given by:

$$Z_t^Y \equiv Z_t X_t^{\frac{\alpha_K}{\alpha_K - 1}} \quad \text{and} \quad Z_t^I \equiv \frac{Z_t^Y}{X_t}.$$

Government spending is assumed to be cointegrated with aggregate output, but with a potentially smoother trend. Letting Z_t^G denote the trend in government spending, its law of motion is:

$$Z_t^G = (Z_{t-1}^G)^{\rho_{zg}} (Z_{t-1}^Y)^{1-\rho_{zg}},$$

where $\rho_{zg} \in [0, 1)$ controls the degree of smoothness in the trend level of government spending.²⁸ To induce stationarity, all model variables are divided by their associated long-run trends. As will be explained in the next section, the equilibrium of the model will be expressed in terms of deviations from these trends.

I assume that the law of motion for each of the nine exogenous processes can be expressed in the form of system (1.5) from section 1.3 in natural logarithms. Specifically, let $z_{\omega,t} \equiv \ln(\omega_t/\omega)$ for $\omega \in \{\zeta_C, \zeta_I, \zeta_Y, \zeta_W, \zeta_P, \zeta_R, g, \mu_Z, \mu_X\}$. The law of motion for each $z_{\omega,t}$ can be written as:

$$\begin{aligned} z_{\omega,t} &= \rho_{z,\omega} z_{\omega,t-1} + a_{\omega,t-1} + e_{\omega,t}^z, & e_{\omega,t}^z &\stackrel{\mathbb{Q}}{\sim} \mathcal{N}(b_{\omega,t-1}, \sigma_{z,\omega}^2) \\ a_{\omega,t} &= \rho_{a,\omega} a_{\omega,t-1} + e_{\omega,t}^a, & e_{\omega,t}^a &\stackrel{\mathbb{Q}}{\sim} \mathcal{N}(0, \sigma_{a,\omega}^2) \\ b_{\omega,t} &= \rho_{b,\omega} b_{\omega,t-1} + e_{\omega,t}^b, & e_{\omega,t}^b &\stackrel{\mathbb{Q}}{\sim} \mathcal{N}(0, \sigma_{b,\omega}^2). \end{aligned} \tag{1.37}$$

1.5 Solution Methods

This section describes the solution method applied to the model from section 1.4. Two features of that model that are out of the ordinary from a computational standpoint are the inclusion of distorted beliefs, and the presence of recursive preferences. The main insight for handling distorted beliefs is to realize that equilibrium policy rules for endogenous variables do not depend on the objective measure — that measure is only relevant for estimation. Therefore the model can be solved under the distorted measure (\mathbb{Q}) and the policy rules can be converted back to the objective measure (\mathbb{P}) for estimation. The main insight related to recursive preferences is that most dynamic equilibrium models like the one in this chapter can be written in a form that makes it possible to use the properties of lognormal random vectors in a self-consistent way to obtain linear-approximate policy rules that do not satisfy certainty equivalence. The presentation is carried out in somewhat general terms, to illustrate that

²⁸Schmitt-Grohé and Uribe (2012) employ this same specification of trend growth in government spending.

these methods can be applied to a wide class of equilibrium models, including those with neither belief distortions nor recursive preferences.

1.5.1 Solving Models with Distorted Beliefs

Beginning with the filtered probability space $(\Omega, \mathfrak{F}, \{\mathfrak{F}_t\}_{t \geq 0}, \mathbb{P})$, consider an economic model of the form:

$$\begin{aligned} y_t &= g(w_t, z_t) \\ w_t &= h(w_{t-1}, z_{t-1}), \end{aligned} \tag{1.38}$$

where z_t is an $n_z \times 1$ vector of all exogenous processes in the model. At the risk of some abuse in notation, this vector includes all exogenous processes, including those that represent advance information or distorted beliefs. The $n_y \times 1$ vector y_t contains all endogenous, non-predetermined variables and the $n_w \times 1$ vector w_t contains all endogenous predetermined variables. I assume that the law of motion for z_t is known and of the form

$$z_t = \lambda^{\mathbb{P}}(z_{t-1}) + e_t, \quad e_t \stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \Sigma_e). \tag{1.39}$$

The vector-valued function $\lambda^{\mathbb{P}} : \mathbb{R}^{n_z} \rightarrow \mathbb{R}^{n_z}$ is known, but the functions $g : \mathbb{R}^{n_w+n_z} \rightarrow \mathbb{R}^{n_y}$ and $h : \mathbb{R}^{n_w+n_z} \rightarrow \mathbb{R}^{n_w}$ are unknown and will be implicitly determined by the set of equilibrium conditions that characterize the equilibrium of the model economy. It is assumed that these conditions can be summarized by a system of non-linear expectational difference equations of the form:

$$E_t^{\mathbb{Q}} [F(y_{t+1}, y_t, w_{t+1}, w_t, z_{t+1}, z_t)] = 1, \tag{1.40}$$

where $F : \mathbb{R}^{2n_y+2n_w+2n_z} \rightarrow \mathbb{R}_+^{n_y+n_w}$ is a known vector-valued function, and the expectation is computed according to the distorted probability measure \mathbb{Q} which may differ from \mathbb{P} . This reflects the fact that the beliefs agents hold at the time they make their decisions may not

coincide with the beliefs implied by their model environment.

While agents do not need to assign probabilities to future states in a way that exactly coincides with \mathbb{P} , they are required to hold beliefs that are not too dissimilar from those implied by that measure. Mathematically, these two measures are assumed to be locally equivalent. That is, at every date $t \geq 0$,

$$\mathbb{Q}(\mathcal{A}) = 0 \Leftrightarrow \mathbb{P}(\mathcal{A}) = 0, \quad \text{for all } \mathcal{A} \in \mathfrak{F}_t.$$

This means that at each point in time, both probability measures agree on which future events will occur with zero probability. Letting \mathbb{P}_t and \mathbb{Q}_t denote the restriction of these probability measures to \mathfrak{F}_t , define the Radon-Nikodým derivative process as

$$M_t \equiv \frac{d\mathbb{Q}_t}{d\mathbb{P}_t} \quad t \geq 0,$$

which is a unique, strictly positive (a.s.) martingale under \mathbb{P} .

This process defines a sequence of conditional distributions for e_{t+1} under \mathbb{Q} . By imposing a particular parametric structure on M_t , it is possible to obtain a closed-form solution for that conditional distribution. Specifically, let $M_0 = 1$ and

$$\begin{aligned} \frac{M_{t+1}}{M_t} = & \left(\frac{|\Psi(z_t)|}{|\Sigma_e|} \right)^{\frac{1}{2}} \exp \left(-\frac{1}{2} [\psi(z_t)' \Psi(z_t)^{-1} \psi(z_t) + e'_{t+1} (\Psi(z_t)^{-1} - \Sigma_e^{-1}) e_{t+1}] \right. \\ & \left. + \psi(z_t) \Psi(z_t)^{-1} e_{t+1} \right) \end{aligned} \quad (1.41)$$

where $\psi : \mathbb{R}^{n_z} \rightarrow \mathbb{R}^{n_z}$ is a known vector function and $\Psi : \mathbb{R}^{n_z} \rightarrow \mathbb{R}^{n_z \times n_z}$ is a known, symmetric positive definite matrix function. It turns out that z_t remains a conditionally normal, exogenous first-order Markov process if and only if M_t is defined in this way. Stating this formally:

Theorem 1. *Under the assumption that $e_t \stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \Sigma_e)$, if $M_t \equiv d\mathbb{Q}_t/d\mathbb{P}_t$ is defined according to the recursion in (1.41), then $e_t | \mathfrak{F}_{t-1} \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(\psi(z_{t-1}), \Psi(z_{t-1}))$ and the law of motion for z_t*

under \mathbb{Q} is:

$$z_t = \lambda^{\mathbb{Q}}(z_{t-1}) + \epsilon_t, \quad \epsilon_t \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(0, \Psi(z_{t-1})),$$

where $\lambda^{\mathbb{Q}}(z_{t-1}) \equiv \lambda^{\mathbb{P}}(z_{t-1}) + \psi(z_{t-1})$ and $\epsilon_t \equiv e_t - \psi(z_{t-1})$. The converse is also true.

In the special case that $\psi(z_t) = 0$ and $\Psi(z_t) = \Sigma_e$, the distorting martingale $M_{t+1} = 1$ at all times, so the conditional distribution of e_t under \mathbb{Q} coincides with its distribution under \mathbb{P} (i.e. there are no belief distortions). In the model from section 1.4, I implicitly assumed $\Psi(z_t) = \Sigma_e$ so that

$$\frac{M_{t+1}}{M_t} = \exp \left(-\frac{1}{2} \psi(z_t)' \Sigma_e^{-1} \psi(z_t) + \psi(z_t) \Sigma_e^{-1} e_{t+1} \right).$$

As pointed out in section 1.3, this implies that distortions in agents' beliefs affect conditional means but not conditional variances, and is desirable for two reasons. First, it ensures that distorted beliefs and advance information are symmetric, because the latter are assumed only to affect conditional means. Second, it ensures that the model remains conditionally homoskedastic under both probability measures.²⁹ Under this parameterization, the law of motion for z_t under \mathbb{Q} is

$$z_t = \lambda^{\mathbb{Q}}(z_{t-1}) + \epsilon_t, \quad \epsilon_t \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(0, \Sigma_e). \quad (1.42)$$

Combining the laws of motion in (1.38) and (1.42) with the equilibrium conditions in (1.40), it follows that the unknown transition rules g and h solve the functional equation:

$$E_t^{\mathbb{Q}} \left[F(g(h(w_t, z_t), \lambda^{\mathbb{Q}}(z_t) + \epsilon_{t+1}), g(w_t, z_t), h(w_t, z_t), w_t, \lambda^{\mathbb{Q}}(z_t) + \epsilon_{t+1}, z_t) \right] = 1. \quad (1.43)$$

²⁹To the extent that homoskedasticity is desirable from an estimation point of view, it is much less troublesome to have conditional heteroskedasticity under the distorted measure than under the objective measure. This is because the model is estimated under the objective measure. Nevertheless, there would be something of a theoretical asymmetry in supposing that subjective volatility is time-varying, while objective volatility is not. And while it would be interesting to investigate situations in which either type of time-variation in volatility is present, that investigation is beyond the scope of this chapter.

While exact solutions for g and h are typically not available, it is possible to approximate these functions using standard numerical methods. Therefore, modeling distorted beliefs in this way does not introduce any computational complexity above what is normally present when solving non-linear dynamic models of this type. The only difference is that once g and h have been obtained (or approximated), estimation is carried out with the dynamics of z_t expressed under the objective measure, as in (1.39).

1.5.2 Linear Approximations and Recursive Preferences

The goal of this subsection is to obtain a linear approximation to the model's equilibrium dynamics in which elements such as recursive preferences affect the determination of endogenous variables.³⁰ The starting point is to compute a log-linear approximation of the function F from (1.43) in terms of the variables w_t , z_t , and ϵ_{t+1} . Letting $f(\tilde{y}, y, \tilde{w}, w, \tilde{z}, z) \equiv \ln(F(\tilde{y}, y, \tilde{w}, w, \tilde{z}, z))$, and using the chain rule together with (1.38) and (1.39),

$$\begin{aligned} & f(g(h(w_t, z_t), \lambda^{\mathbb{Q}}(z_t) + \epsilon_{t+1}), g(w_t, z_t), h(w_t, z_t), w_t, \lambda^{\mathbb{Q}}(z_t) + \epsilon_{t+1}, z_t) \\ & \approx f + [f_{\tilde{y}}g_w h_w + f_y g_w + f_{\tilde{w}}h_w + f_w](w_t - \bar{w}) \\ & + [f_{\tilde{y}}(g_w + g_z \lambda_z^{\mathbb{Q}}) + f_y g_z + f_{\tilde{w}}h_z + f_{\tilde{z}}\lambda_z^{\mathbb{Q}} + f_z](z_t - \bar{z}) \\ & + [f_{\tilde{y}}g_z + f_{\tilde{z}}]\epsilon_{t+1}, \end{aligned}$$

where $f_x \equiv \frac{\partial f}{\partial x'}$, and it is understood that all constant terms are evaluated at the steady state levels $w_t = \bar{w}$, $z_t = \bar{z}$, and $\epsilon_{t+1} = 0$.³¹ Substituting this approximation into (1.43), and

³⁰To further clarify, the main aim is really to obtain an approximation where *risk premia* have non-zero effects on the approximate dynamics. Recursive preferences simply provide an additional degree of freedom to allow those premia to better fit the data.

³¹This subsection retains the distinction between the distorted and historical probability measures. For models in which those two measures coincide, replace \mathbb{Q} by \mathbb{P} throughout.

taking natural logarithms of both sides implies that:

$$\begin{aligned}
0 = & f + [f_{\tilde{y}}g_w h_w + f_y g_w + f_{\tilde{w}}h_w + f_w](w_t - \bar{w}) \\
& + [f_{\tilde{y}}(g_w + g_z \lambda_z^{\mathbb{Q}}) + f_y g_z + f_{\tilde{w}}h_z + f_{\tilde{z}}\lambda_z^{\mathbb{Q}} + f_z](z_t - \bar{z}) \\
& + \frac{1}{2} \text{diag}([f_{\tilde{y}}g_z + f_{\tilde{z}}]\Sigma_e[f_{\tilde{y}}g_z + f_{\tilde{z}}]') .
\end{aligned} \tag{1.44}$$

The last term in this expression makes use of the fact that ϵ_{t+1} is conditionally normal under \mathbb{Q} , together with the properties of lognormally distributed random vectors. For any square matrix A , The operator $\text{diag}(A)$ collects the diagonal elements of A into a column vector.

Because expression (1.44) must hold for any realization of w_t and z_t , it must be the case that each of the constant terms and two coefficient terms are equal to zero. Consider first the term multiplying w_t :

$$f_{\tilde{y}}g_w h_w + f_y g_w + f_{\tilde{w}}h_w + f_w = 0. \tag{1.45}$$

As long as the number of stable generalized eigenvalues of the matrices $-[f_{\tilde{w}}, f_{\tilde{y}}]$ and $[f_w, f_y]$ equals the number of endogenous predetermined variables n_w , it is possible to solve this polynomial matrix equation for g_w and h_w .³²

Next, the coefficient on z_t is used to determine g_z and h_z :

$$f_{\tilde{y}}(g_w + g_z \lambda_z^{\mathbb{Q}}) + f_y g_z + f_{\tilde{w}}h_z + f_{\tilde{z}}\lambda_z^{\mathbb{Q}} + f_z = 0, \tag{1.46}$$

where in this expression, g_w and h_w are known (in addition to $\lambda_z^{\mathbb{Q}}$). Vectorizing this system, and using the fact that $\text{vec}(ABC) = (C' \otimes A)\text{vec}(B)$,

$$[(I_{n_z} \otimes f_y) + ((\lambda_z^{\mathbb{Q}})' \otimes f_{\tilde{y}}), \quad (I_{n_z} \otimes [f_{\tilde{w}} + f_{\tilde{y}}g_w])] \begin{bmatrix} \text{vec}(g_z) \\ \text{vec}(h_z) \end{bmatrix} = -\text{vec}(f_z + f_{\tilde{z}}\lambda_z^{\mathbb{Q}}).$$

³²For details on the Schur decomposition method employed in this chapter, see [Klein \(2000\)](#).

This linear system determines g_z and h_z .³³

Lastly, the constant terms \bar{w} and $\bar{y} \equiv g(\bar{w}, \bar{z})$ in (1.46), which are arguments of the known derivative matrices, satisfy the system of $n_y + n_w$ nonlinear equations:

$$f + \frac{1}{2} \text{diag}([f_{\bar{y}}g_z + f_{\bar{z}}]\Sigma_e[f_{\bar{y}}g_z + f_{\bar{z}}]') = 0. \quad (1.47)$$

The usual result that $f(\bar{y}, \bar{y}, \bar{w}, \bar{w}, \bar{z}, \bar{z}) = 0$ no longer holds, and it is generally not possible to obtain closed form solutions for all these constant terms. Instead, these equations can be solved numerically. One consideration that greatly simplifies the task of solving this system is that it is possible to show that the rows of $f_{\bar{y}}$ and $f_{\bar{z}}$ corresponding to some of the equilibrium conditions is zero. Using the fact that for these conditions, the corresponding element of f must still equal zero, it is often possible to derive closed-form expressions for a number of steady state relations that hold almost surely under \mathbb{Q} (and therefore also under \mathbb{P}). This reduces the number of unknown values that must be determined numerically.

How is this method helpful in cases with recursive preferences? Primarily because it delivers a linear solution where the matrices g_w , g_z , h_w , h_z , and the steady-state values \bar{y} , \bar{w} are functions of the parameter γ , which adjusts the household's continuation value for risk. The risk adjustment affects the steady state of the model through the second term on the left-hand side of (1.47), which then affects the transition dynamics through (1.45) and (1.46). If the second term on the left-hand side of (1.47) is omitted, then this solution method is equivalent to a first-order perturbation approximation of the model around its non-stochastic steady state, as in [Schmitt-Grohé and Uribe \(2004\)](#) for example. Formally,

Theorem 2. *The linear-approximate dynamics derived in this subsection are equivalent to a first-order perturbation approximation of the model's dynamics around its non-stochastic steady state if and only if the steady-state values \bar{y} and \bar{w} satisfy the $n_y + n_w$ system of*

³³An alternative approach to solving (1.46) writes the system in the form of a generalized Sylvester equation. For details on this approach, see [Gomme and Klein \(2011\)](#).

equations

$$f(\bar{y}, \bar{y}, \bar{w}, \bar{w}, \bar{z}, \bar{z}) = 0,$$

taking \bar{z} as given.

As an example of how this solution method works, consider the hypothetical asset-pricing exercise of obtaining a linear solution to the real risk-free rate based on the following equilibrium condition:

$$1 = E_t^{\mathbb{Q}} \left[\beta \left(\frac{C_{t+1}}{C_t} \right)^{-1/\xi} \left(\frac{V_{t+1}}{E_t^{\mathbb{Q}}[V_{t+1}^{1-\gamma}]^{1/(1-\gamma)}} \right)^{1/\xi-\gamma} R_t^f \right],$$

where $\Delta c_{t+1} \equiv \ln(C_{t+1}/C_t)$ and $vc_{t+1} \equiv \ln(V_{t+1}/C_{t+1})$ are conditionally linear *exogenous* processes. A standard first order perturbation approximation of this expression around the non-stochastic steady state implies that

$$r_t^f = -\ln(\beta) + \frac{1}{\xi} E_t^{\mathbb{Q}}[\Delta c_{t+1}].$$

Note that γ does not appear in this expression, so no amount of data on r_t^f , Δc_t , or vc_t would identify that parameter when this approximation is used.

In order to apply the approximation method of this subsection, it is first necessary to rewrite the nonlinear condition in the form of (1.40). Introducing the auxiliary variable $X_t \equiv E_t^{\mathbb{Q}}[V_{t+1}^{1-\gamma}]^{1/(1-\gamma)}$, the system is

$$1 = E_t^{\mathbb{Q}} \left[\frac{\exp \left(\ln(\beta) - \frac{1}{\xi} \Delta c_{t+1} + \left(\frac{1}{\xi} - \gamma \right) [vc_{t+1} + \Delta c_{t+1} - x_t] \right)}{\exp \left((1-\gamma)[vc_{t+1} + \Delta c_{t+1} - x_t] \right)} \right],$$

where $x_t \equiv \ln(X_t)$. In this case the function F is log-linear, so no approximations are necessary. Using the properties of lognormal random vectors and substituting x_t out of the

result implies that

$$r_t^f = -\ln(\beta) - \frac{1}{\xi} E_t^{\mathbb{Q}}[\Delta c_{t+1}] - \frac{1}{2} \gamma^2 \sigma_{\Delta c}^2 - \frac{1}{2} \left(\frac{1}{\xi} - \gamma \right) \left[(1 - \gamma) \sigma_{\Delta c}^2 + \left(1 - \frac{1}{\xi} \right) \sigma_{vc}^2 + \sigma_{\Delta c, vc} \right],$$

where $\sigma_{\Delta c}^2$ is the conditional variance of Δc_t , σ_{vc}^2 is the conditional variance of vc_t , and $\sigma_{\Delta c, vc}$ is their conditional covariance. From this expression, it is clear that γ influences the equilibrium level of $E^{\mathbb{Q}}[r_t^f]$. This expression for r_t^f is exact because Δc_t and vc_t are assumed to be linear processes for the purpose of this example. If that were not the case, the solution procedure entails replacing their exact laws of motion with linear approximations.

1.6 Empirical Analysis

In this section, the model from section 1.4 is confronted with data. A subset of its parameters are calibrated to conventional values from related studies, and the remaining parameters are estimated using likelihood-based methods. The primary focus is on the importance of belief-driven fluctuations, and the relative importance of advance information and distorted beliefs for different series. As a brief summary of the main findings, the estimates point to a large role for belief-driven fluctuations. In particular, belief distortions are most important for explaining stock returns and consumption growth, and advance information is most important for explaining inflation and investment growth.

1.6.1 Estimation Method

The first step in the estimation procedure is to describe the mapping from a collection of observable time series to their model counterparts. The approximation method in section 1.5 produces a linear system characterizing the model's dynamics. The state vector in the economy is given by $x_t \equiv (w'_t, z'_t)'$. An $n_o \times 1$ observation vector y_t^o is related to this state

vector by the equation

$$y_t^o = \bar{y}^o + Ax_t^* + \eta_u u_t, \quad u_t \stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \Sigma_u), \quad (1.48)$$

where $x_t^* \equiv x_t - \bar{x}$. The $n_u \times 1$ vector u_t contains independent and identically distributed measurement errors, which affect y_t^o according to the $n_o \times n_u$ selection matrix η_u . The matrix A mapping the observation vector to the latent states can be written as

$$A \equiv S_o \begin{bmatrix} g_x \\ h_x \\ g_x h_x \\ h_x h_x \end{bmatrix}, \quad \text{where} \quad g_x \equiv [g_w, g_z], \quad h_x \equiv \begin{bmatrix} h_w & h_z \\ 0 & \lambda_z^{\mathbb{Q}} \end{bmatrix},$$

and S_o is an $n_o \times 2(n_y + n_x)$ selection matrix. The constant term \bar{y}^o is defined as $\bar{y}^o \equiv A\bar{x}$. The presence of the last two blocks in A permit expectations of future model variables to be included in the observation vector. To complete the state-space representation of the model in terms of y_t^o and x_t^* , the dynamics of x_t^* must be expressed under the objective measure:

$$x_t^* = \begin{bmatrix} h_w & h_z \\ 0 & \lambda_z^{\mathbb{P}} \end{bmatrix} x_{t-1}^* + \begin{bmatrix} 0 \\ I_{n_e} \end{bmatrix} e_t, \quad e_t \stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \Sigma_e). \quad (1.49)$$

The transition and covariance matrices in (1.48) and (1.49) are functions of the underlying parameter vector Θ . The goal of estimation is to approximate the posterior distribution of Θ conditional on the complete sample of data $\{y_t^o\}_{t=1}^T$. By Bayes' rule, the posterior combines information from the likelihood function and the prior distribution:

$$p(\Theta | \{y_t^o\}_{t=1}^T) \propto p(\{y_t^o\}_{t=1}^T | \Theta) p(\Theta).$$

Because the system is linear, the likelihood function can be evaluated using the Kalman filter.

This is a main advantage of the solution method employed in this chapter. An estimate of the mode of the posterior distribution is first obtained by numerical optimization, and then the entire distribution is simulated using Markov Chain Monte Carlo (MCMC) methods.³⁴

1.6.2 Data

The model is estimated using quarterly U.S. data from 1954:Q3 – 2009:Q1. This sample excludes the period after 2009:Q1 when the nominal interest rate was at its lower bound. The fourteen data series used in estimation, along with the model concepts associated with these series, are listed in Table 1.1.

Model Concept	Description
$\Delta \ln(Y_t^d)$	Real per capita GDP growth
$\Delta \ln(C_t)$	Real per capita consumption growth
$\Delta \ln(X_t I_t)$	Real per capita investment growth
$\Delta \ln(H_t)$	Growth in per capita hours worked
$\Delta \ln(W_t)$	Real wage growth
$\ln(\Pi_t)$	Inflation
$\ln(R_t^e)$	Real stock market return
$\ln(R_t)$	Nominal risk-free return
$E_t^Q[\Delta \ln(Y_{t+1}^d)]$	Expected per capita real GDP growth
$E_t^Q[\ln(\Pi_{t+1})]$	Expected inflation
$E_t^Q[\ln(R_{t+1})]$	Expected nominal risk-free return
$\Delta \ln(G_t)$	Real per capita government spending growth
$\Delta \ln(\text{TFP}_t)$	Total factor productivity growth
$\Delta \ln(X_t)$	Growth in relative price of investment

Table 1.1: Data. The variable $\text{TFP}_t \equiv \zeta_t^Y Z_t^{1-\alpha_K}$ denotes (capacity-adjusted) total factor productivity, and Δ is the linear first-difference operator. The sample period is 1954:Q3 – 2009:Q1. A more detailed description of the data and sources is included in Appendix A.1.

Listing them here for convenience as well, they are: growth in real per capita GDP, consumption, investment, and government spending; growth in per capita hours worked, real wages, total factor productivity, and the relative price of investment; inflation, real stock returns, the nominal risk-free rate; and finally expected growth in per capita real

³⁴See An and Schorfheide (2007) for a detailed description of this approach.

GDP, inflation, and nominal risk-free returns. Three series do not extend all the way back to 1954:Q3. Expected output and inflation only extend to 1968:Q4 and the expected nominal risk-free rate only extends back to 1981:Q3. In these cases, observations before the available start date are treated as missing observations. All remaining variables contain observations over the entire sample.

Three different groups of these variables have been included to help discipline the estimation. First, the real stock market return and three expectational series contain the most direct evidence on agents' beliefs about the future. They are important for plausibly identifying the relative importance of advance information and distorted beliefs, which is the main focus of the empirical analysis. Second, the growth rates of real government consumption, total factor productivity, and the relative price of investment are fundamentals in the model. Including observations on these variables is helpful to ensure that the postulated fundamental processes (and implicitly the beliefs about them) behave consistently with the available observations. Third, wage growth, inflation, and the nominal interest rate discipline the nominal rigidities in the model: the imperfect price-setting structure in labor and retail markets, and the monetary authority's feedback rule.

1.6.3 Calibrated Parameters and Priors

Some parameters contained in Θ are calibrated. These parameters, their fixed values, and a brief description are listed in Table 1.2. The marginal rate of time preference β is fixed at 0.995, which corresponds to a real interest rate of about 2% in the non-stochastic steady state. A value of one for the elasticity of intertemporal substitution is standard within the business-cycle literature. The share of non-labor inputs (capital and land) in production is assumed to be 0.325, and the labor share is 0.675. As in [Jaimovich and Rebelo \(2009\)](#) and [Schmitt-Grohé and Uribe \(2012\)](#), the degree of decreasing returns to scale in production is fixed at 10 percent, which implies that the capital share is 0.225. In the depreciation function, $\delta(u) = \delta_0 + \delta_1(u - 1) + \delta_2/2(u - 1)^2$, the parameter δ_1 is chosen to ensure that

capacity utilization, u , equals one in the steady state. In the steady state, capital depreciates at a rate of 10 percent per year, which corresponds to $\delta_0 = 0.025$.

Parameter	Value	Description
β	0.995	Marginal rate of time preference
ξ	1	Elasticity of intertemporal substitution
α_K	0.225	Capital share in production
α_H	0.675	Labor share in production
δ_0	0.025	Steady-state depreciation rate
u	1	Steady-state utilization rate
μ^Y	1.0044	Steady-state gross growth rate of per capita GDP
μ^X	0.9968	Steady state gross growth rate of investment price
G/Y^d	0.2103	Steady state share of government purchases in GDP
h	0.2	Steady-state level of hours
$\eta_W/(\eta_W - 1)$	1.15	Steady-state wage markup
$\eta_P/(\eta_P - 1)$	1.15	Steady-state price markup
φ_R	0	Persistence in nominal interest rate rule
Π	1.0086	Steady-state inflation rate

Table 1.2: Calibrated parameters. The time unit is one quarter.

The steady-state growth rates μ^Y and μ^X , the share of government purchases G/Y^d , and the steady-state inflation rate Π are chosen to match their average levels over the sample period. The parameter ϑ , which scales the disutility of labor in the period utility kernel, is chosen to ensure that the steady-state fraction of time spent working (h) is equal to 20 percent. The steady-state wage and price markups are fixed at 15 percent, which is a value in line with [Justiniano et al. \(2011\)](#). The interest rate feedback rule of the monetary authority does not feature any endogenous persistence, $\varphi_R = 0$. A common alternative is to assume an “inertial” policy rule ($\varphi_R > 0$), with independent and identically distributed errors. The specification chosen here allows all exogenous disturbances in the model to be treated symmetrically. Moreover, [Rudebusch \(2002, 2006\)](#) and [Carrillo et al. \(2007\)](#) provide empirical evidence that a policy rule with serially correlated disturbances and little internal persistence may be more consistent with the dynamic behavior of interest rates.

The remaining parameters in Θ that are not calibrated are assigned prior distributions. These all fall closely in line with the existing literature, and are reported in [Table 1.3](#). The

prior for risk aversion, γ is assumed to be a (generalized) gamma distribution with a mean of five and standard deviation of two. Without habit persistence or labor supply, γ represents the risk aversion parameter of the household over atemporal wealth gambles. A value of five for this parameter would in that case lie in the center of the conventional range for this parameter in the asset pricing literature since [Mehra and Prescott \(1985\)](#). The leverage parameter λ is assigned a gamma distribution with a mean of 2 and standard deviation of 0.75. According to this prior, aggregate stock returns are twice as volatile as the growth rate in the value of productive firms. This is consistent with the amount of financial leverage measured by [Rauh and Sufi \(2012\)](#), and is consistent with the calibration of [Croce \(2014\)](#).

For the exogenous processes, the persistence parameters of contemporaneous fundamental disturbances are assigned (generalized) beta priors with a mean of 0.6 and standard deviation of 0.2. The persistence of permanent technology and the relative price of investment are assigned a lower mean of 0.2, to reflect the view that these series (absent any advance information) are close to being independent and identically distributed across time. To allow for the possibility that these parameters may be exactly zero, or even somewhat negative, the priors are bounded below by negative one rather than zero. The remaining persistence parameters associated with advance information and distorted beliefs are assigned beta priors centered at zero.

The prior distribution for the standard deviation of each contemporaneous disturbance is inverse gamma with a mean of 0.5 and a standard deviation of 1. The standard deviations associated with advance information and distorted beliefs are instead centered at 0.1. In the event that all persistence parameters were fixed at zero, these priors imply that contemporaneous disturbances are responsible for about 93 percent of the variation in fundamentals under the \mathbb{Q} -measure, and about 96 percent under the \mathbb{P} -measure.

Parameter	Description	Dist.	Mean	S.D.	L.B.	U.B.
γ	risk aversion	\mathcal{G}	5	2	1	$+\infty$
λ	equity leverage	\mathcal{G}	2	0.75	1	$+\infty$
ν	wealth elasticity of labor	\mathcal{U}	0.5	$1/(2\sqrt{3})$	0	1
ϕ	habit persistence	\mathcal{B}	0.5	0.1	0	1
ψ	labor supply elasticity	\mathcal{G}	2	0.75	1	$+\infty$
κ	investment adj costs	\mathcal{G}	4	1	0	$+\infty$
δ_2/δ_1	utilization cost elasticity	\mathcal{IG}	0.5	0.5	0	$+\infty$
θ_W	wage no reset probability	\mathcal{B}	0.66	0.1	0	1
θ_P	price no reset probability	\mathcal{B}	0.66	0.1	0	1
χ_W	wage indexation	\mathcal{B}	0.5	0.15	0	1
χ_P	price indexation	\mathcal{B}	0.5	0.15	0	1
φ_Π	feedback rule inflation	\mathcal{G}	1.7	0.3	1	$+\infty$
φ_{Y_d}	feedback rule output growth	\mathcal{G}	0.125	0.05	0	$+\infty$
ρ_{ZG}	smooth trend govt spending	\mathcal{B}	0.6	0.2	-1	1
$\rho_{z,i}$	persistence current fundamentals	\mathcal{B}	0.6	0.2	-1	1
ρ_{z,μ_Z}	persistence perm tech growth	\mathcal{B}	0.2	0.2	-1	1
ρ_{z,μ_X}	persistence rel price of inv	\mathcal{B}	0.2	0.2	-1	1
$\rho_{a,i}$	persistence advance information	\mathcal{B}	0	0.2	-1	1
$\rho_{b,i}$	persistence belief distortions	\mathcal{B}	0	0.2	-1	1
$\sigma_{z,i}$	std dev current fundamentals	\mathcal{IG}	0.5	1	0	$+\infty$
$\sigma_{z,i}$	std dev advance information	\mathcal{IG}	0.1	1	0	$+\infty$
$\sigma_{z,i}$	std dev belief distortions	\mathcal{IG}	0.1	1	0	$+\infty$
σ_{u,Y^d}	measurement error GDP growth	\mathcal{B}	0.1	0.05	0	0.2923

Table 1.3: Prior distributions. The symbol \mathcal{U} denotes a uniform distribution, \mathcal{B} is a generalized beta distribution, \mathcal{G} is a generalized gamma distribution, and \mathcal{IG} is an inverse-gamma distribution. The generalized distributions are parameterized by a location parameter, scale parameter, lower bound, and upper bound. The columns titled L.B. and U.B. represent the lower and upper bounds defining the support of these distributions. The upper bound for the measurement error is equal to $\sqrt{0.10} \times \text{std}(\Delta \ln(Y_t^d))$, so that the variance of this error is at most ten percent of the observed variance in GDP growth.

1.6.4 Estimation Results

Table 1.4 displays estimates of the structural parameters in the model, and Tables 1.5 and 1.6 display estimates of the parameters that govern the exogenous processes in the model. For comparison, the prior means are included in the table as well. The two asset-pricing parameters in the model are γ , which controls the risk preferences of the household, and λ , which controls the degree of leverage in equity returns relative to the growth in firm value. The former is estimated to be close to ten: somewhat higher than the prior mean and near the upper limit of the conventional range (0, 10) typically imposed on this parameter in the case without labor supply or habit persistence. The leverage parameter is slightly higher than its prior mean at 2.7, but still consistent with existing studies.

Parameter	Prior	Posterior				Description
	Mean	Mean	Mode	10th	90th	
γ	5	10.50	10.37	9.18	11.86	risk aversion
λ	2	2.66	2.70	2.22	3.05	equity leverage
ν	0.5	0.06	0.06	0.05	0.07	wealth elasticity of labor
ϕ	0.5	0.68	0.68	0.65	0.70	habit persistence
ψ	2	1.20	1.20	1.15	1.25	labor supply elasticity
κ	4	13.9	13.5	12.27	15.6	investment adj costs
δ_2/δ_1	0.5	0.13	0.13	0.11	0.15	utilization cost elasticity
θ_W	0.66	0.03	0.03	0.02	0.04	wage no reset probability
θ_P	0.66	0.08	0.08	0.06	0.10	price no reset probability
χ_W	0.5	0.50	0.51	0.31	0.68	wage indexation
χ_P	0.5	0.26	0.20	0.13	0.41	price indexation
φ_Π	1.7	4.99	4.77	4.36	5.67	feedback rule inflation
φ_{Y_d}	0.125	0.23	0.22	0.15	0.31	feedback rule output growth
ρ_{ZG}	0.6	0.63	0.69	0.40	0.82	smooth trend govt spending

Table 1.4: Estimated structural parameters. The mean, 10th, and 90th percentiles come from 100,000 draws from a Markov chain approximating the posterior distribution of the parameters. These draws are taken from 10 Markov chains of length 10,000 after removing a burn-in sample of 1,000 draws.

The remaining estimates of preference and technology parameters are consistent with existing literature. For example, the parameter ν , which controls the degree of wealth effects on labor supply, is small at 0.06. This is not quite as small as the value estimated by [Schmitt-](#)

Grohé and Uribe (2012), who find a posterior median estimate of $\nu = 0.003$. Nevertheless, this estimate still confirms their finding that wealth effects are small, and that preferences are closer to the type of Greenwood et al. (1988).

The degree of nominal rigidity in the model is estimated to be small. The probabilities that wages and prices are reset each period are 0.97 and 0.92, respectively. These imply that the average time between optimal readjustments in labor and product markets is close to three months. By contrast, Justiniano et al. (2011) find an average duration around one year for optimal wages and prices. The degree of price indexation is also small, at 26 percent. Some studies have estimated smaller values: Justiniano et al. (2011) find only ten percent indexation in both markets. But, others have worked with higher values: in their empirical model, Altig et al. (2011) calibrate both indexation parameters to one. The estimated feedback rule contains a strong response to increases in inflation, and a smaller response to output growth.

To gauge the ability of the model to fit the data, Table 1.7 presents some selected moments from the model and the data. While the estimation procedure attempts to match the entire autocovariance function of the vector of observable variables, this comparison can be helpful for understanding the model's limitations. Overall, the fit is comparable to other related studies. According to the model, consumption growth is less volatile than output growth, investment growth is much more volatile than output growth, equity returns are much more volatile than the risk-free rate, and there is an ex-post equity premium of about four percent, which almost exactly matches the 4.05% equity premium observed in the data.

Nevertheless, the model does fall short along several dimensions. For one thing, it overstates the standard deviation of almost every variable. It also overstates the first-order autocorrelation of output and consumption by about an order of two. At the same time, the model understates the autocorrelation of hours, as well as its correlation with output. In terms of nominal variables, the model overstates the degree of negative correlation between the nominal risk-free rate and output growth.

Parameter	Prior	Posterior				Type	Fundamental
	Mean	Mean	Mode	10th	90th		
ρ_{z,ζ_C}	0.6	0.75	0.75	0.62	0.86	C	preference
ρ_{z,ζ_I}	0.6	0.03	0.02	-0.05	0.11	C	inv. specific
ρ_{z,ζ_Y}	0.6	0.98	0.98	0.98	0.99	C	temp. productivity
ρ_{z,ζ_W}	0.6	0.99	0.99	0.99	1.00	C	union marginal cost
ρ_{z,ζ_P}	0.6	1.00	1.00	1.00	1.00	C	retailer marginal cost
ρ_{z,ζ_R}	0.6	0.18	0.18	0.14	0.23	C	monetary policy
$\rho_{z,g}$	0.6	0.96	0.97	0.95	0.98	C	gov. spending
ρ_{z,μ_Z}	0.2	0.69	0.72	0.58	0.78	C	permanent productivity
ρ_{z,μ_X}	0.2	0.46	0.47	0.40	0.50	C	rel. price of inv.
ρ_{a,ζ_C}	0	0.05	0.03	-0.20	0.31	A	preference
ρ_{a,ζ_I}	0	0.73	0.73	0.69	0.77	A	inv. specific
ρ_{a,ζ_Y}	0	0.29	0.28	0.08	0.49	A	temp. productivity
ρ_{a,ζ_W}	0	0.98	0.98	0.98	0.99	A	union marginal cost
ρ_{a,ζ_P}	0	0.29	0.28	0.08	0.50	A	retailer marginal cost
ρ_{a,ζ_R}	0	0.90	0.90	0.86	0.92	A	monetary policy
$\rho_{a,g}$	0	0.01	0.00	-0.24	0.27	A	gov. spending
ρ_{a,μ_Z}	0	0.06	0.04	-0.16	0.27	A	permanent productivity
ρ_{a,μ_X}	0	-0.03	-0.01	-0.30	0.24	A	rel. price of inv.
ρ_{b,ζ_C}	0	0.04	0.04	-0.22	0.30	B	preference
ρ_{b,ζ_I}	0	-0.00	-0.00	-0.27	0.26	B	inv. specific
ρ_{b,ζ_Y}	0	0.29	0.28	0.08	0.50	B	temp. productivity
ρ_{b,ζ_W}	0	0.20	0.19	-0.02	0.42	B	union marginal cost
ρ_{b,ζ_P}	0	0.99	0.99	0.99	1.00	B	retailer marginal cost
ρ_{b,ζ_R}	0	0.85	0.85	0.79	0.89	B	monetary policy
$\rho_{b,g}$	0	0.72	0.73	0.66	0.78	B	gov. spending
ρ_{b,μ_Z}	0	0.55	0.52	0.47	0.65	B	permanent productivity
ρ_{b,μ_X}	0	0.81	0.81	0.79	0.84	B	rel. price of inv.

Table 1.5: Estimated persistence parameters of exogenous processes. “Type” denotes the type of exogenous process, either: current fundamentals (C), advance information (A), or belief distortions (B). The column titled “Fundamental” provides a brief description of the fundamental process associated with each type of disturbance. The mean, 10th, and 90th percentiles come from 100,000 draws from a Markov chain approximating the posterior distribution of the parameters. These draws are taken from 10 Markov chains of length 10,000 after removing a burn-in sample of 1,000 draws.

Parameter	Prior	Posterior				Type	Fundamental
	Mean	Mean	Mode	10th	90th		
σ_{z,ζ_C}	0.5	0.41	0.41	0.22	0.60	C	preference
σ_{z,ζ_I}	0.5	21.22	20.91	18.60	24.02	C	inv. specific
σ_{z,ζ_Y}	0.5	0.84	0.83	0.78	0.89	C	temp. productivity
σ_{z,ζ_W}	0.5	0.30	0.23	0.18	0.46	C	union marginal cost
σ_{z,ζ_P}	0.5	0.75	0.75	0.71	0.79	C	retailer marginal cost
σ_{z,ζ_R}	0.5	1.10	1.04	0.94	1.27	C	monetary policy
$\sigma_{z,g}$	0.5	1.05	1.05	0.98	1.11	C	gov. spending
σ_{z,μ_Z}	0.5	0.14	0.13	0.12	0.17	C	permanent productivity
σ_{z,μ_X}	0.5	0.30	0.31	0.28	0.33	C	rel. price of inv.
σ_{a,ζ_C}	0.1	0.05	0.04	0.03	0.07	A	preference
σ_{a,ζ_I}	0.1	6.91	6.89	5.65	8.24	A	inv. specific
σ_{a,ζ_Y}	0.1	0.03	0.03	0.02	0.04	A	temp. productivity
σ_{a,ζ_W}	0.1	0.10	0.10	0.08	0.12	A	union marginal cost
σ_{a,ζ_P}	0.1	0.03	0.03	0.02	0.04	A	retailer marginal cost
σ_{a,ζ_R}	0.1	0.50	0.47	0.41	0.60	A	monetary policy
$\sigma_{a,g}$	0.1	0.07	0.05	0.04	0.13	A	gov. spending
σ_{a,μ_Z}	0.1	0.04	0.04	0.03	0.05	A	permanent productivity
σ_{a,μ_X}	0.1	0.07	0.05	0.04	0.13	A	rel. price of inv.
σ_{b,ζ_C}	0.1	0.05	0.04	0.03	0.07	B	preference
σ_{b,ζ_I}	0.1	0.09	0.05	0.04	0.16	B	inv. specific
σ_{b,ζ_Y}	0.1	0.03	0.03	0.02	0.04	B	temp. productivity
σ_{b,ζ_W}	0.1	0.03	0.03	0.03	0.04	B	union marginal cost
σ_{b,ζ_P}	0.1	0.04	0.04	0.03	0.05	B	retailer marginal cost
σ_{b,ζ_R}	0.1	0.44	0.41	0.35	0.53	B	monetary policy
$\sigma_{b,g}$	0.1	1.55	1.52	1.34	1.77	B	gov. spending
σ_{b,μ_Z}	0.1	0.15	0.14	0.12	0.17	B	permanent productivity
σ_{b,μ_X}	0.1	0.57	0.56	0.51	0.62	B	rel. price of inv.
σ_{u,Y_d}	0.1	0.29	0.29	0.29	0.29	M	output growth

Table 1.6: Estimated volatility parameters of exogenous processes. “Type” denotes the type of exogenous process, either: current fundamentals (C), advance information (A), belief distortions (B), or measurement error (M). The column titled “Fundamental” provides a brief description of the fundamental process associated with each type of disturbance. The mean, 10th, and 90th percentiles come from 100,000 draws from a Markov chain approximating the posterior distribution of the parameters. These draws are taken from 10 Markov chains of length 10,000 after removing a burn-in sample of 1,000 draws.

Variable	Mean	Std Dev	AC(1)	Corr GDP
$\Delta \ln(Y_t^d)$	1.76 (1.76)	2.21 (1.85)	0.60 (0.34)	1.00 (1.00)
$\Delta \ln(C_t)$	1.76 (2.01)	1.67 (1.09)	0.65 (0.29)	0.68 (0.55)
$\Delta \ln(X_t I_t)$	1.76 (2.01)	6.19 (4.56)	0.62 (0.44)	0.86 (0.77)
$\Delta \ln(H_t)$	0.00 (-0.15)	2.63 (1.75)	0.24 (0.61)	0.53 (0.73)
$\Delta \ln(W_t)$	1.76 (1.64)	2.78 (1.28)	-0.02 (0.01)	0.30 (0.14)
$\Delta \ln(\Pi_t)$	3.43 (3.43)	0.92 (1.16)	0.74 (0.86)	-0.24 (-0.25)
$\ln(R_t^e)$	4.66 (5.82)	23.66 (17.04)	-0.04 (0.09)	0.17 (0.13)
$\ln(R_t)$	4.17 (5.20)	2.08 (1.43)	0.87 (0.96)	-0.31 (-0.10)
$E_t^{\mathbb{Q}}[\Delta \ln(Y_{t+1}^d)]$	1.76 (1.15)	1.82 (0.92)	0.82 (0.70)	0.66 (0.53)
$E_t^{\mathbb{Q}}[\ln(\Pi_{t+1})]$	3.43 (3.82)	0.86 (0.99)	0.89 (0.96)	-0.19 (-0.15)
$E_t^{\mathbb{Q}}[\ln(R_{t+1})]$	4.17 (5.41)	1.84 (1.36)	0.88 (0.97)	-0.30 (0.10)

Table 1.7: Selected moments in the model and data. Moments from the data are listed in parentheses below their model-implied counterparts. Means and standard deviations are computed at the quarterly frequency and then annualized by multiplying by 4 (means) or 2 (standard deviations). Units are in percentage points.

1.6.5 Decomposing Fluctuations

Table 1.8 displays the share of the unconditional variance of each variable used in estimation that can be attributed to current fundamental disturbances, advance information, and distorted beliefs. The first main finding is that for most endogenous variables (the first eleven rows), beliefs play an important role. For example, over 40 percent of the fluctuations in output growth and 70 percent of the fluctuations in consumption growth are due either to advanced information or distorted beliefs. On average across all endogenous variables, two thirds of the variation can be attributed to one of these two sources. The second finding is that both advance information and distorted beliefs play an economically significant role. Neither type of disturbance completely crowds out the other across the board. In particular, both are equally important for explaining output growth and hours worked.

Advance information and distorted beliefs are not equally important for all series however. Advance information is most important for inflation (and expected inflation), at over 50%. The fact that inflation reflects disturbances that were correctly anticipated is consistent with the idea that the actions of the Federal Reserve are largely known in advance. Indeed, anticipated monetary policy disturbances are responsible for over half of the variation in inflation (53%). This can be seen in Table 1.9, which displays the contributions from the selected exogenous processes that are most important for explaining the observed endogenous variables. Advance information is also particularly important for real investment (44%). Nearly all of this anticipated component comes from advance information related to transitory investment-specific technological change. Advance information about changes in investment-specific technology also explain almost all of the anticipated component of GDP growth.

On the other hand, distorted beliefs are much more important for explaining stock returns and consumption growth. This result is noteworthy because those two series are the main focus of many modern consumption-based asset-pricing models. It is consistent with the idea that the prices of financial assets are less tied to fundamental developments than are

Variable	Current	Advance	Distortions
$\Delta \ln(Y_t^d)$	0.51 (0.47,0.55)	0.27 (0.22,0.31)	0.22 (0.19,0.25)
$\Delta \ln(C_t)$	0.23 (0.19,0.28)	0.14 (0.12,0.17)	0.62 (0.58,0.67)
$\Delta \ln(X_t I_t)$	0.50 (0.45,0.56)	0.44 (0.37,0.50)	0.06 (0.05,0.07)
$\Delta \ln(H_t)$	0.41 (0.38,0.44)	0.24 (0.21,0.28)	0.34 (0.32,0.37)
$\Delta \ln(W_t)$	0.61 (0.58,0.64)	0.05 (0.04,0.06)	0.34 (0.31,0.38)
$\ln(\Pi_t)$	0.33 (0.28,0.38)	0.55 (0.48,0.62)	0.12 (0.09,0.15)
$\ln(R_t^e)$	0.16 (0.13,0.20)	0.11 (0.07,0.16)	0.72 (0.65,0.80)
$\ln(R_t)$	0.29 (0.25,0.33)	0.20 (0.18,0.23)	0.50 (0.46,0.55)
$E_t^Q[\Delta \ln(Y_{t+1}^d)]$	0.36 (0.32,0.39)	0.31 (0.26,0.36)	0.34 (0.29,0.39)
$E_t^Q[\ln(\Pi_{t+1})]$	0.08 (0.06,0.11)	0.61 (0.53,0.69)	0.31 (0.23,0.38)
$E_t^Q[\ln(R_{t+1})]$	0.27 (0.24,0.31)	0.21 (0.18,0.25)	0.51 (0.47,0.56)
$\Delta \ln(G_t)$	0.99 (0.98,1.00)	0.01 (0.00,0.02)	0 (-)
$\Delta \ln(\text{TFP}_t)$	1.00 (0.99,1.00)	0.00 (0.00,0.01)	0 (-)
$\Delta \ln(X_t)$	0.93 (0.84,0.99)	0.07 (0.01,0.16)	0 (-)

Table 1.8: Variance decomposition. Each column contains the contribution from a different type of disturbance. “Current” denotes current fundamentals, “Advance” denotes advance information, and “Distortions” denotes belief distortions. Decompositions are computed based on a random sample of 1000 parameters from the posterior distribution. The main entry represents the mean value, and the numbers in parentheses below are the 10th and 90th percentiles, respectively.

real investment, hiring, and production. What are these distorted beliefs about? In the case of stock returns, distorted beliefs about future marginal costs in product markets represent the largest factor (51%), while distorted beliefs about the relative price of investment are the second largest (19%). Distorted beliefs about the relative price of investment are also important for consumption growth (36%), but future marginal costs play less of a role there (2%). Instead, distortions related to future government purchases are more important (17%).

Among the contemporaneous innovations to current fundamentals, those to investment-specific technology and temporary productivity are the most important for explaining output (36%), investment (42%), and hours (33%). Innovations to product-market marginal costs play an important role in explaining wage growth (30%). These are presented in the first three columns of Table 1.9. These findings are consistent with those in the existing literature. In particular, the importance of investment-specific technology for explaining these variables is also documented by [Justiniano et al. \(2011\)](#) and [Schmitt-Grohé and Uribe \(2012\)](#). Also consistent with the second of these studies is the finding that the anticipated component of investment-specific technology plays a large role. Indeed, for output, investment, and hours, the anticipated innovation in this fundamental process is even slightly more important than the unanticipated innovation. Finally, a new contribution to the discussion of investment-specific disturbances is the finding in Table 1.9 that distorted beliefs concerning the relative price of investment explain a substantial fraction of the variation in output (10%), consumption (36%), hours (20%), wages (30%), stock returns (19%) and nominal interest rates (12%). Indeed, distorted beliefs about this process are the single most important source of distortions for the real economy.

In contrast to existing studies (e.g. [Smets and Wouters, 2007](#)), exogenous variation in labor market marginal costs does not represent a particularly large source of fluctuations in the data. Current innovations in this process make up less than one percent of the variation in all endogenous variables. The advance information component is larger, but still explains less than ten percent of the variation in each of the endogenous variables. [Schmitt-Grohé](#)

Variable	Current			Advance			Distortions		
	ζ_I	ζ_Y	ζ_P	ζ_I	ζ_W	ζ_R	ζ_P	g	μ_X
$\Delta \ln(Y_t^d)$	0.14	0.22	0.05	0.21	0.05	0.00	0.00	0.07	0.10
$\Delta \ln(C_t)$	0.01	0.11	0.02	0.08	0.06	0.00	0.02	0.17	0.36
$\Delta \ln(X_t I_t)$	0.21	0.21	0.06	0.40	0.03	0.00	0.01	0.02	0.01
$\Delta \ln(H_t)$	0.10	0.23	0.03	0.15	0.09	0.00	0.01	0.10	0.20
$\Delta \ln(W_t)$	0.00	0.29	0.30	0.02	0.03	0.00	0.01	0.02	0.32
$\ln(\Pi_t)$	0.00	0.07	0.01	0.01	0.01	0.53	0.06	0.01	0.03
$\ln(R_t^e)$	0.00	0.09	0.06	0.02	0.09	0.00	0.51	0.00	0.19
$\ln(R_t)$	0.00	0.22	0.04	0.03	0.06	0.11	0.27	0.01	0.12
$E_t^{\mathbb{Q}}[\Delta \ln(Y_{t+1}^d)]$	0.01	0.22	0.06	0.23	0.07	0.00	0.01	0.20	0.07
$E_t^{\mathbb{Q}}[\ln(\Pi_{t+1})]$	0.00	0.06	0.01	0.01	0.01	0.59	0.05	0.01	0.02
$E_t^{\mathbb{Q}}[\ln(R_{t+1})]$	0.00	0.21	0.03	0.03	0.06	0.11	0.26	0.01	0.12

Table 1.9: Variance decomposition for selected exogenous processes. Each column contains the contribution from a different exogenous process. The first set of three columns contain the contributions from the three current disturbances that are most important for explaining fluctuations in the eleven endogenous variables: investment-specific technology ζ_I , temporary neutral productivity ζ_Y , and product market marginal costs ζ_P . The second set of three columns contain the contributions from the three exogenous processes with the largest advance information component: investment-specific technology ζ_I , labor market marginal costs ζ_W , and monetary policy ζ_R . The last set of three columns contain the contributions from the three exogenous processes with the largest belief distortion component: product market marginal costs ζ_P , government spending g , and the relative price of investment μ_X . Each entry in the table represents the mean value based on a random sample of 1000 parameters from the posterior distribution (the same random sample used in Table 1.8).

Variable	Current	Advance	Distortions
$\Delta \ln(Y_t^d)$	0.92	0.04	0.03
$\Delta \ln(C_t)$	0.91	0.05	0.04
$\Delta \ln(X_t I_t)$	0.94	0.04	0.02
$\Delta \ln(H_t)$	0.79	0.12	0.09
$\Delta \ln(W_t)$	0.94	0.04	0.02
$\ln(\Pi_t)$	0.94	0.05	0.01
$\ln(R_t^e)$	0.94	0.02	0.04
$\ln(R_t)$	0.95	0.03	0.01
$E_t^Q[\Delta \ln(Y_{t+1}^d)]$	0.92	0.05	0.03
$E_t^Q[\ln(\Pi_{t+1})]$	0.92	0.06	0.02
$E_t^Q[\ln(R_{t+1})]$	0.91	0.06	0.04
$\Delta \ln(G_t)$	0.96	0.04	0
$\Delta \ln(\text{TFP}_t)$	0.96	0.04	0
$\Delta \ln(X_t)$	0.96	0.04	0

Table 1.10: Variance decomposition at the prior mean. Each column contains the contribution from a different type of disturbance. The decomposition is computed at the prior mean values listed in Table 1.3. Column headings are the same as those in Table 1.8.

and Uribe (2012), for example, find that the anticipated component of this process explains 67 percent of the variation in hours, and between 15 and 20 percent of the variation in output and consumption. This is a particularly appealing result given the amount of criticism leveled against labor supply shocks of this type (e.g. by Chari et al., 2009; Shimer, 2009).

Last but not least, what about the survey forecasts? For the most part, Table 1.8 indicates that the importance of current fundamentals, advance information, and distorted beliefs for these variables largely mirror their realized counterparts. Expected inflation largely reflects advance information as does realized inflation. Belief distortions are important for expected nominal risk-free interest rates, as they are for realized interest rates. However, distorted beliefs are somewhat more important for expected output than they are for realized output (34% vs 22%). On average, 62 percent of the variation in the three forecast series is attributable to either current or (correctly) anticipated changes in fundamentals.

As a point of comparison, Table 1.10 displays the same variance decomposition as Table 1.8, but with the parameters set to their prior means. In every case disturbances to current fundamentals play the largest role. Stock returns and consumption are not as strongly

affected by belief distortions, and inflation does not have such a large anticipated component. While these observations are informative, it is also important to keep in mind that the prior distributions were chosen to reflect uncertainty surrounding the parameters themselves, not about the model-implied variance decompositions.

1.7 Robustness

This section explores the robustness of the findings from section 1.6 to four independent perturbations of interest. First, I remove stock returns from the set of observable variables. Second, I add classical measurement error to the three survey forecasts. Third I remove belief distortions from the analysis altogether to see what the data have to say when the only alternative to current innovations is advance information. Fourth, I increase the degree of wage and price rigidity in the model by fixing the parameters θ_W and θ_P to their prior means. In each case, results are only reported for the posterior mode. In the baseline case for example, Tables 1.4-1.6 show that the posterior mode does a good job of capturing the central tendency of the underlying parameter draws.

1.7.1 No Stock Returns

Stock returns have been included in the analysis so far for three main reasons. First, because a large literature in financial economics suggests that the stock market should contain useful information about the market's expectations of future developments in the economy. Second, because conclusions about the relative importance of advance information and distorted beliefs may be (and in fact are) different for macroeconomic and financial variables. And third, because it is of independent interest whether a medium-scale model of the type developed in this chapter is capable of explaining the joint dynamics of financial and macroeconomic variables — especially those aspects that have been difficult for existing models, such as the equity premium and excess volatility of stock returns. However, because standard practice

in the business cycle literature is not to force the model to be consistent with data on stock returns, this subsection considers how things change if stock returns are omitted from the set of observable variables.

Variable	Current	Advance	Distortions
$\Delta \ln(Y_t^d)$	0.51	0.27	0.22
$\Delta \ln(C_t)$	0.22	0.12	0.67
$\Delta \ln(X_t I_t)$	0.48	0.47	0.06
$\Delta \ln(H_t)$	0.49	0.18	0.33
$\Delta \ln(W_t)$	0.66	0.03	0.31
$\ln(\Pi_t)$	0.34	0.58	0.07
$\ln(R_t^e)$	0.35	0.09	0.57
$\ln(R_t)$	0.40	0.20	0.40
$E_t^{\mathbb{Q}}[\Delta \ln(Y_{t+1}^d)]$	0.38	0.30	0.32
$E_t^{\mathbb{Q}}[\ln(\Pi_{t+1})]$	0.10	0.63	0.28
$E_t^{\mathbb{Q}}[\ln(R_{t+1})]$	0.37	0.20	0.42
$\Delta \ln(G_t)$	1.00	0.00	0
$\Delta \ln(\text{TFP}_t)$	1.00	0.00	0
$\Delta \ln(X_t)$	0.98	0.02	0

Table 1.11: Variance decomposition without stock returns. The decomposition is computed at the posterior mode.

Table 1.11 displays the variance decomposition based on the estimates without stock returns. The decomposition of stock returns are still included in this table even though they are no longer an observable series, for comparison with Table 1.8. For the most part the results are similar to the baseline case. The most notable deviations are that without data on stock returns, the importance of belief distortions for stock returns and the nominal risk-free rate both fall by about ten percent. Instead, these variations are now attributed to innovations in current fundamentals. In particular, the fraction that was previously attributed to distorted beliefs about marginal costs in product markets is now attributed to current innovations in transitory productivity. Nevertheless it still remains true that belief distortions are most important for stock returns and consumption, and that advance information is most important for inflation and investment.

1.7.2 Measurement Errors in Survey Forecasts

A main emphasis of this chapter has been on the importance of using direct data on subjective beliefs to determine the relative importance of advance information and distorted beliefs about fundamentals. However, it is possible to entertain the view that while such direct data is valuable in principle, the forecasts used in practice for the empirical analysis of this chapter may be subject to errors in measurement. To allow for that possibility, I introduce classical measurement errors on each of the three survey forecasts and re-estimate the variances of those errors along with the other parameters in the model. The standard deviations of these errors are assigned inverse gamma priors with a standard deviation of one; the means are chosen so that at the prior mean, the measurement error variance is one percent of its observable counterpart.

Variable	Current	Advance	Distortions
$\Delta \ln(Y_t^d)$	0.52	0.28	0.20
$\Delta \ln(C_t)$	0.24	0.13	0.63
$\Delta \ln(X_t I_t)$	0.50	0.45	0.05
$\Delta \ln(H_t)$	0.42	0.23	0.35
$\Delta \ln(W_t)$	0.62	0.04	0.34
$\ln(\Pi_t)$	0.42	0.42	0.15
$\ln(R_t^e)$	0.13	0.08	0.79
$\ln(R_t)$	0.33	0.18	0.50
$E_t^Q[\Delta \ln(Y_{t+1}^d)]$	0.39	0.31	0.30
$E_t^Q[\ln(\Pi_{t+1})]$	0.16	0.56	0.28
$E_t^Q[\ln(R_{t+1})]$	0.30	0.19	0.51
$\Delta \ln(G_t)$	1.00	0.00	0
$\Delta \ln(\text{TFP}_t)$	1.00	0.00	0
$\Delta \ln(X_t)$	0.98	0.02	0

Table 1.12: Variance decomposition with classical measurement error in survey forecasts. The decomposition is computed at the posterior mode.

Table 1.12 displays the variance decomposition of all observable variables after the measurement errors have been included. The results are essentially unchanged from the baseline case. Evidently the data favors the model's structural driving processes over measurement errors for explaining the observed variation in survey forecasts. For the forecasts of output

growth and the nominal risk-free rate, the measurement error explains less than one percent of the series' total variation. This number is slightly higher for expected inflation, at about six percent, but still small compared to the other disturbances in the model. The results are therefore robust to having measurement errors of this type in the survey forecasts.

1.7.3 No Distorted Beliefs

The focus of this chapter is on the relative importance of advance information *and* distorted beliefs for explaining macroeconomic and financial data. Nevertheless, it would be interesting to know what different conclusions might have been reached if distorted beliefs were excluded from the start as an alternative hypothesis. To investigate this possibility, I eliminate the distorted belief component of each fundamental processes so that the subjective probabilities coincide with those implied by the model. I then re-estimate the remaining parameters using the same priors listed in Table 1.3.

The variance decomposition based on this exercise are reported in Table 1.13. The main result is that without distorted beliefs, advance information is responsible for about half of the variation in almost every endogenous variable, and more than half for expected output growth (64%) and expected inflation (83%). Moreover, the predictable component is now larger for all three of the exogenous processes. Only 11% in the case of total factor productivity, but nearly forty percent for government spending. These findings are roughly consistent with those of [Schmitt-Grohé and Uribe \(2012\)](#).

However, it is advance information about transitory neutral technology and investment-specific technology rather than marginal costs in the labor market that are responsible for the bulk of the anticipated components of output, consumption, hours, and wages.³⁵ Inflation is still largely driven by advance information about the exogenous component of monetary policy, and stock returns by advance information about marginal costs in product markets. Current innovations to marginal costs in the labor market do play a larger role in explaining

³⁵The decomposition by exogenous process are not shown; the salient points are reported here in the text.

Variable	Current	Advance	Distortions
$\Delta \ln(Y_t^d)$	0.54	0.46	0
$\Delta \ln(C_t)$	0.57	0.43	0
$\Delta \ln(X_t I_t)$	0.51	0.49	0
$\Delta \ln(H_t)$	0.56	0.44	0
$\Delta \ln(W_t)$	0.59	0.41	0
$\ln(\Pi_t)$	0.45	0.55	0
$\ln(R_t^e)$	0.60	0.40	0
$\ln(R_t)$	0.50	0.50	0
$E_t^{\mathbb{Q}}[\Delta \ln(Y_{t+1}^d)]$	0.36	0.64	0
$E_t^{\mathbb{Q}}[\ln(\Pi_{t+1})]$	0.17	0.83	0
$E_t^{\mathbb{Q}}[\ln(R_{t+1})]$	0.51	0.49	0
$\Delta \ln(G_t)$	0.62	0.38	0
$\Delta \ln(\text{TFP}_t)$	0.89	0.11	0
$\Delta \ln(X_t)$	0.73	0.27	0

Table 1.13: Variance decomposition without distorted beliefs. The decomposition is computed at the posterior mode.

hours at 17 percent, compared to less than one percent in the baseline case with distorted beliefs.

1.7.4 More Nominal Rigidity

One unusual feature of the results in Table 1.4 relative to related business cycle models with nominal rigidities is that the degree of wage and price rigidity is estimated to be low: wages and prices may be reset optimally almost every quarter. The purpose of this subsection is to examine the extent to which the main findings in section 1.6 are robust to imposing a greater degree of nominal price and wage rigidity. To do so, I fix both θ_P and θ_W at their prior mean value of 0.66 and re-estimate the model. These values imply that the average time between optimal readjustments in labor and product markets is around nine months.

Table 1.14 displays the variance decomposition results. In this case, the importance of advance information and distorted beliefs generally declines for all variables. For example, half of consumption growth is explained by distorted beliefs (down from 62% in the baseline case), and only two percent is explained by advance information (down from 14% in the base-

line case). On average over all eleven endogenous variables, these two belief-related sources explain around 40 percent of the observed variation. It is still true that belief distortions are most important for explaining stock returns and consumption growth, although they are now relatively more important for consumption growth.

Variable	Current	Advance	Distortions
$\Delta \ln(Y_t^d)$	0.67	0.07	0.26
$\Delta \ln(C_t)$	0.48	0.02	0.50
$\Delta \ln(X_t I_t)$	0.88	0.10	0.02
$\Delta \ln(H_t)$	0.65	0.10	0.25
$\Delta \ln(W_t)$	0.43	0.28	0.29
$\ln(\Pi_t)$	0.60	0.23	0.17
$\ln(R_t^e)$	0.32	0.37	0.31
$\ln(R_t)$	0.67	0.16	0.17
$E_t^Q[\Delta \ln(Y_{t+1}^d)]$	0.49	0.09	0.43
$E_t^Q[\ln(\Pi_{t+1})]$	0.57	0.24	0.19
$E_t^Q[\ln(R_{t+1})]$	0.62	0.18	0.21
$\Delta \ln(G_t)$	1.00	0.00	0
$\Delta \ln(\text{TFP}_t)$	1.00	0.00	0
$\Delta \ln(X_t)$	0.98	0.02	0

Table 1.14: Variance decomposition with more nominal rigidity. The two parameters governing the frequency of (optimal) price adjustment are set to $\theta_W = \theta_P = 0.66$. The decomposition is computed at the posterior mode.

Which current fundamental has become a more plausible explanation in the presence of larger nominal rigidity? Labor market marginal costs. Current innovations to this exogenous process are now responsible for 20 percent of the observed variation in output, 15 percent in consumption, 15 percent in hours worked, 13 percent in real wages, and 10 percent of inflation. By comparison, this innovation explained less than one percent of the variation in each endogenous variable according to the baseline results.

1.8 Two Historical Episodes

So far, the relative importance of advance information and distorted beliefs has been discussed in terms of the forecast error variance decomposition implied by the estimated model.

Another way to discuss their role in driving fluctuations is to consider the roles they have played in specific historical episodes. This section presents estimates of the sources of fluctuations in the real economy and the stock market during two historical episodes: the Dot-Com boom and bust of the mid-late 1990s and the recent Great Recession. These episodes are given special attention because the first is one in which developments in the stock market took center stage, and because the second is the largest downturn in the real economy since the Great Depression. Therefore, they provide interesting case studies of both macroeconomic and financial fluctuations. Before describing the results, I briefly describe the procedure used to construct the historical decompositions.

1.8.1 Construction of Historical Decompositions

The state-space model in (1.48) and (1.49) implies a specific conditional density for the exogenous disturbance vector $\varepsilon_t \equiv (u'_t, e'_t)'$. Given the set of observable variables and a value for the parameter vector, it is possible to compute estimates of these disturbances using the Kalman smoother. By performing counterfactual experiments where some disturbances are “shut off” (set to zero) and others are not, it is possible to isolate the contribution of each shock to the different observable variables. The details of this procedure are presented in Algorithm 1.³⁶

1.8.2 The Dot-Com Boom and Bust

In the 1990’s U.S. stock prices increased by almost a factor of five — the largest increase during any decade in the nation’s history. The “Dot-Com Boom,” as this period has come to be known, was marked by unprecedented growth in information technology, computers, and internet-based companies. Following this period of growth was a sharp and rapid bust. From its peak at the turn of the century, the stock market lost nearly half its value over the

³⁶For a discussion of how to compute smoothed estimates of the exogenous disturbances, which is the first step of the algorithm, see [Durbin and Koopman \(2012\)](#).

Algorithm 1 Historical decomposition of fluctuations

- (1) Use the Kalman smoother to estimate the sequence of exogenous disturbances $\{\hat{\varepsilon}_t\}_{t=1}^T$ and the initial state \hat{x}_0^* based on the posterior mean parameter estimate $\hat{\Theta}$, where

$$\hat{x}_0^* \equiv E^{\mathbb{P}}[x_0^* | \{y_t^o\}_{t=1}^T, \hat{\Theta}] \quad \text{and} \quad \hat{\varepsilon}_t \equiv E^{\mathbb{P}}[\varepsilon_t | \{y_t^o\}_{t=1}^T, \hat{\Theta}].$$

- (2) For each disturbance $i = 1, \dots, n_\varepsilon$ (the i -th element of ε_t),
- (a) Construct a new sequence $\{\hat{\varepsilon}_t(i)\}_{t=1}^T$ by replacing all elements but those of the i -th series in $\hat{\varepsilon}_t$ by zeros.
 - (b) Beginning with \hat{x}_0^* , use the state dynamics in equation (1.49) and the observation equation (1.48) to generate a sequence of counterfactual observations $\{\hat{y}_t^o(i)\}_{t=1}^T$ based on $\{\hat{\varepsilon}_t(i)\}_{t=1}^T$.
- (3) The value $\hat{y}_t^o(i)$ represents the estimated contribution of disturbance i at date t to the observed variation in y_t^o .
-

next year and a half. What caused these dramatic fluctuations? The estimated model of this chapter provides a precise answer to that question.

Figure 1.2 shows the sources of fluctuations in per capita output growth and aggregate stock returns over the ten-year period from 1992 to 2002. The relative importance of current fundamentals, advance information, and distorted beliefs during this episode are generally consistent with the “unconditional” importance of these three sources reported in Table 1.8. Most of the fluctuations in output growth are the result of developments in current fundamentals, followed by advance information and distorted beliefs. By contrast, most of the fluctuations in stock returns are the result of belief distortions; however, current fundamentals also play a non-negligible role.

What type of picture does the model paint for the stock market boom and bust during these years? For most of the period from 1995 to 1999, the stock market was buffeted by a series of positive innovations in current fundamentals — specifically, expansionary innovations to current transitory productivity and product market marginal costs ($e_{\zeta_Y,t}^z$ and $e_{\zeta_P,t}^z$ in the notation of section 1.4). These capture the positive effects of technological development during the period. After this sequence of innovations, the stock market was

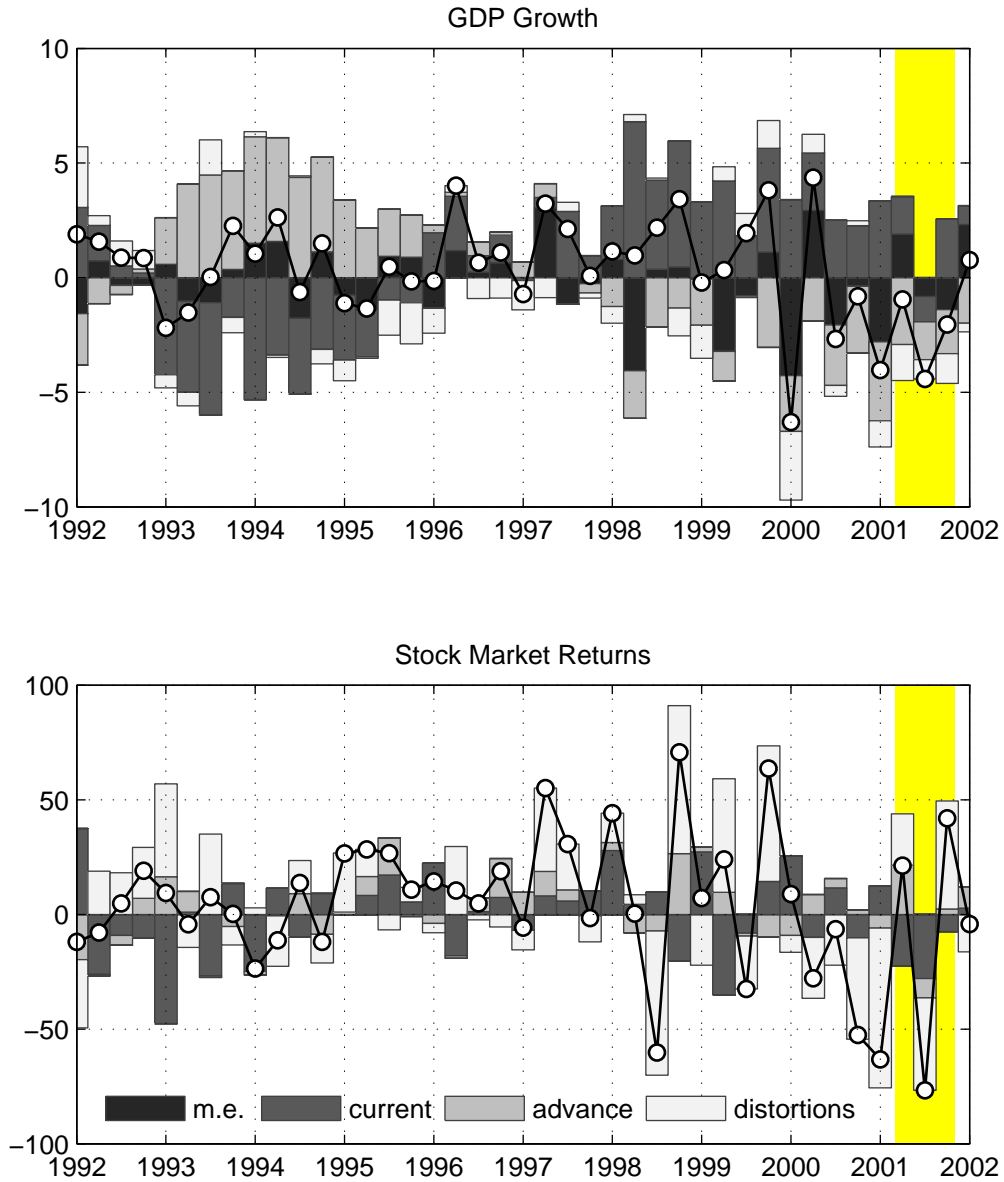


Figure 1.2: Dot-Com Boom and Bust (1992-2002). The top panel is the annualized growth rate of real per capita GDP and the bottom panel is the annualized ex-post return on the aggregate stock market. Each quarterly observation (white circle) is decomposed into one of four sources: measurement error (m.e.), current fundamentals (current), advance information about future fundamentals (advance), or belief distortions about future fundamentals (distortions). These add the contributions from all exogenous processes of that respective category. Contributions from individual disturbances are discussed in the main text. The yellow shaded region is a recessionary period according to the NBER Business Cycle Dating Committee.

affected by three quarters of particularly excessive optimism: in the last quarter of 1998, the second quarter of 1999, and the last quarter of 1999. The peak of the boom was in the first quarter of 2000. The bust was fueled by a sequence of large negative belief distortions about marginal costs in product markets ($e_{\zeta_P,t}^b$). These distortions appear to have had something like a grain of truth to them. In each quarter that beliefs about future marginal costs were excessively pessimistic, *actual* (current) marginal costs were also negative.

In the real economy, the positive innovations in current productivity that contributed to the boom in the stock market also led to a boom in GDP growth. In fact, the persistent effects of these innovations on output lasted until even after the peak of the stock market in 2000. At that time, a sequence of adverse information about future fundamentals led to negative output growth, ultimately culminating in the 2001 recession. The belief distortions about future marginal costs that drove the stock market bust were less important for the real economy. That is because those innovations generate offsetting effects on output growth. The expectation of persistently higher marginal costs in product markets lead households to expect higher current and future prices, which reduces consumption. The fact that households are averse to late resolution of uncertainty implies that they discount risky projects more heavily as a result. Equilibrium in asset markets requires the expected return on investment to increase, leading to an increase in investment by productive firms. This increase in investment offsets the initial decrease in consumption leading to only a small response in aggregate output.

Taken together, the picture is one of technological innovations leading to a boom in both the real economy and the stock market. A wave of optimism at the end of the boom helped to propel the stock market into the stratosphere. A sudden wave of pessimism at the turn of the century, accompanied by some actual negative contemporaneous developments in firms' marginal costs, drove the crash. Because underlying productivity in the economy was still fairly strong, however, aggregate output did not suffer that much. The subsequent recession was only a mild one by historical standards.

1.8.3 The Great Recession

The real economy only experienced a mild downturn after the Dot-Com boom of the 1990s, but it would experience a severe collapse just a few years later. In the fourth quarter of 2008 alone, per capita real GDP fell by almost ten percent on an annualized basis — the largest single quarter decline in over half a century. At the same time, the stock market collapse around this time was just as severe as the one that occurred in 2000. Figure 1.3 displays the joint behavior of per capita GDP growth and aggregate stock returns from 2002 to 2009, together with the model-implied account of the sources of fluctuations in these series.

According to the model, output began falling in 2006 and 2007 due to a series of negative realizations in current transitory neutral productivity ($e_{\zeta_Y,t}^z$) from 2006 to the beginning of 2008. The largest of these occurred in the first quarter of 2008. The cataclysmic drop in output growth that followed in the fourth quarter of 2008 was primarily caused by a large negative innovation to current transitory investment-specific productivity ($e_{\zeta_I,t}^z$). Investment-specific productivity captures the economy’s ability to transform current savings into future capital inputs. Justiniano et al. (2011) argue that this type of disturbance is likely to proxy for more fundamental disturbances to the functioning of the financial sector. That interpretation is consistent with the idea that the bankruptcy of the large U.S. investment bank Lehman Brothers in the fourth quarter of 2008 severely impeded the functioning of the financial sector, leading to a sharp drop in output growth.

Belief distortions also played an important role. During the downturn, the real economy experienced a wave of pessimism concerning future permanent neutral productivity ($\mu_{Z,t}$) and product market marginal costs (ζ_t^P). This pessimism had two effects. First, it amplified the drop in output growth beyond what would have occurred with current fundamental disturbances alone. Second, it served as the primary cause for the collapse of the stock market. This is seen from the second panel of Figure 1.3. Even though current disturbances in investment specific productivity had large effects in the real economy (output, consumption, investment, and hours worked), they had almost no impact on the stock market.

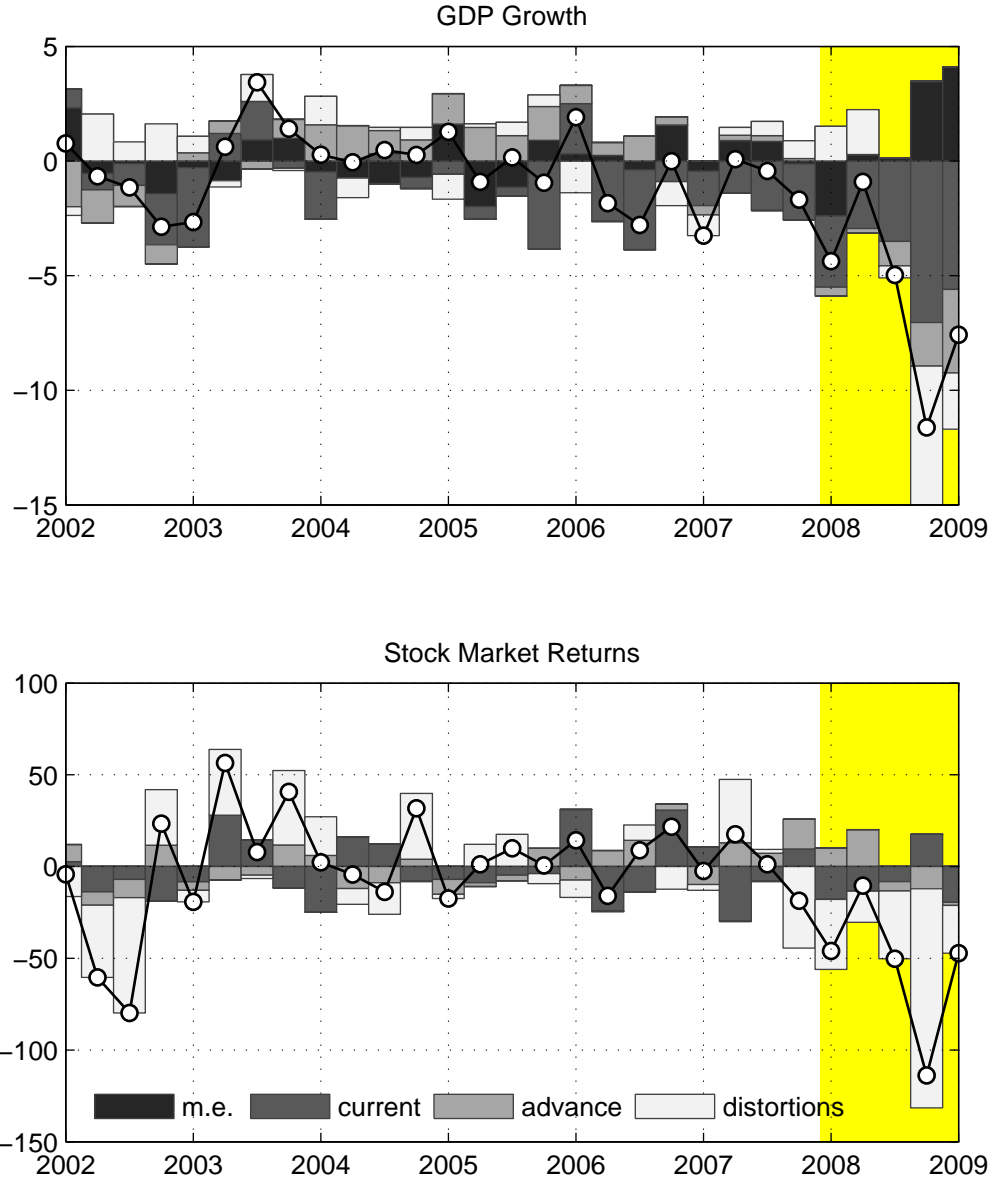


Figure 1.3: Run-up to the Great Recession (2002-2009). The top panel is the annualized growth rate of real per capita GDP and the bottom panel is the annualized ex-post return on the aggregate stock market. Each quarterly observation (white circle) is decomposed into one of four sources: measurement error (m.e.), current fundamentals (current) advance information about future fundamentals (advance), or belief distortions about future fundamentals (distortions). These add the contributions from all exogenous processes of that respective category. Contributions from individual disturbances are discussed in the main text. The yellow shaded region is a recessionary period according to the NBER Business Cycle Dating Committee.

Turning to the period before the recession, both output growth and stock returns remained fairly close to their respective trends. In 2003, a few positive innovations in transitory neutral productivity contributed to a small upturn in both series. A conspicuous feature of the decomposition during this period is that there is no strong upward pressure coming from monetary policy shocks. Mechanically, this is because the estimated degree of nominal rigidity in the model is low. Although nominal interest rates were low during most of this period, that was a result of the monetary authority responding to low inflation. In fact, policymakers in the U.S. at that time were quite explicit about their concern over low inflation. For example, in his 2003 testimony before the U.S. Congress, Federal Reserve Chairman Alan Greenspan justified low policy rates as an appropriate response to low inflation, and even expressed concern over the possibility of a deflationary episode.³⁷

While this description of the sources of fluctuations behind the Great Recession may be plausible, it is necessarily incomplete. In particular, there are two important things to keep in mind. First, the model in this chapter does not feature a housing sector, so it has nothing to say about the boom in house prices that occurred during this period. Second, while the model is capable of capturing disruptions in the financial sector in a reduced-form way — through innovations to transitory investment-specific productivity — it would certainly be desirable to incorporate a more microfounded model of financial frictions. With those considerations in mind, the conclusion that comes through from this decomposition analysis is that the Great Recession was primarily driven by negative contemporaneous developments in transitory neutral and investment-specific productivity, together with a wave of pessimism about future permanent neutral productivity and marginal costs in product markets.

³⁷His exact words were:

“...[W]e face new challenges in maintaining price stability, specifically to prevent inflation from falling too low. This is one reason the FOMC has adopted a quite accommodative stance of policy. A very low inflation rate increases the risk that an adverse shock to the economy would be more difficult to counter effectively. Indeed, there is an especially pernicious, albeit remote, scenario in which inflation turns negative against a backdrop of weak aggregate demand, engendering a corrosive deflationary spiral.” [Greenspan \(2003\)](#)

1.9 Conclusion

This chapter has investigated the role of advance information and distorted beliefs in explaining the observed fluctuations in a range of macroeconomic and financial time series, and has presented three main findings. First, advance information and distorted beliefs are important. About two-thirds of the variation in endogenous variables can be attributed to one of these two sources. Second, both sources are separately important; neither one completely swamps the other. Third, these two sources important in different ways for different variables. Advance information is most important for explaining inflation and investment, while distorted beliefs are most important for explaining stock returns and consumption. They are roughly equally important for explaining output growth and hours worked.

In arriving at these results, this chapter has built and estimated a new medium-scale DSGE model that incorporates recursive utility preferences, internal habit formation, capital accumulation, and nominal rigidities, and which is capable of delivering an empirically adequate description of a set of key U.S. macroeconomic and financial data. It has also introduced a flexible way of incorporating belief distortions into a class of optimizing equilibrium models that is useful for empirical work. Specifically, it has proposed to directly parameterize the Radon-Nikodým process that links agents' subjective beliefs to the objective ones implied by their model environment. Finally, it has developed a new internally consistent framework for obtaining a linear approximation to models that naturally incorporates important “higher-order” information. This framework is shown to be particularly useful for models in which agents are not indifferent about the intertemporal resolution of uncertainty.

The findings in this chapter also suggest several fruitful areas of future research. For one, given the importance of advance information and distorted beliefs, what are the precise mechanisms responsible for generating these underlying sources of fluctuations? One of the main reasons for developing structural models like the one in this chapter is to perform counterfactual experiments that can be helpful for characterizing the effects of different policy actions. But to perform experiments of that type, it is important to know how the econ-

omy’s technology for producing advance information or distorted beliefs would change under different policy regimes. The exogenous specifications employed in this chapter represent a first approximation that is helpful for arriving at positive results, but not for performing counterfactual policy analysis.

A second interesting question involves the role of uncertainty. Several recent papers have emphasized the importance of time-varying uncertainty for explaining macroeconomic phenomena.³⁸ This chapter has abstracted from those types of disturbances in order to focus more sharply on different sources of fluctuations in beliefs about conditional means. However, essentially all of the arguments and methods presented here would directly apply to cases when agents have advance information and distorted beliefs about conditional variances. In fact, it turns out that the solution procedures developed in section 1.5 can be directly applied to obtain conditionally linear approximations of models with time-varying uncertainty. This may seem either surprising or unhelpful, given that conventional wisdom says that linear approximations are not capable of capturing the effects of time-varying volatility. But the main intuition for why this is not necessarily the case is as follows: if time-varying uncertainty is of “first-order” importance, it should be possible to construct a linear (i.e. first-order) approximation which captures those effects. The solution method in this chapter can be extended to provide one internally consistent way of doing just that.

³⁸For example, Bloom (2009), Fernández-Villaverde et al. (2011), and Bloom et al. (2012).

Chapter 2

Measuring Uncertainty¹

2.1 Introduction

How important is time-varying economic uncertainty and what role does it play in macroeconomic fluctuations? A large and growing body of literature has concerned itself with this question.² At a general level, uncertainty is typically defined as the conditional volatility of a disturbance that is unforecastable from the perspective of economic agents. In partial equilibrium settings, increases in uncertainty can depress hiring, investment, or consumption if agents are subject to fixed costs or partial irreversibilities (a “real options” effect), if agents are risk averse (a “precautionary savings” effect), or if financial constraints tighten in response to higher uncertainty (a “financial frictions” effect). In general equilibrium settings, many of these mechanisms continue to imply a role for time-varying uncertainty, although some may also require additional frictions to generate the same effects.

A challenge in empirically examining the behavior of uncertainty, and its relation to macroeconomic activity, is that no objective measure of uncertainty exists. So far, the em-

¹This chapter was co-authored with Sydney C. Ludvigson and Serena Ng. It has been published as: “Measuring Uncertainty,” *American Economic Review*, 105(3):1177-1216.

²See for example, Bloom (2009); Arellano et al. (2012) Bloom et al. (2012); Bachmann et al. (2013); Gilchrist et al. (2014); Schaal (2011); Bachmann and Bayer (2011); Baker and Bloom (2013); Baker et al. (2013); Basu and Bundick (2012); Knotek and Khan (2011); Fernández-Villaverde et al. (2011); Leduc and Liu (2014); Nakamura et al. (2012); Orlik and Veldkamp (2014).

pirical literature has relied primarily on proxies or indicators of uncertainty, such as the implied or realized volatility of stock market returns, the cross-sectional dispersion of firm profits, stock returns, or productivity, the cross-sectional dispersion of subjective (survey-based) forecasts, or the appearance of certain “uncertainty-related” key words in news publications. While most of these measures have the advantage of being directly observable, their adequacy as proxies for uncertainty depends on how strongly they are correlated with this latent stochastic process.

Unfortunately, the conditions under which common proxies are likely to be tightly linked to the typical theoretical notion of uncertainty may be quite special. For example, stock market volatility can change over time even if there is no change in uncertainty about economic fundamentals, if leverage changes, or if movements in risk aversion or sentiment are important drivers of asset market fluctuations. Cross-sectional dispersion in individual stock returns can fluctuate without any change in uncertainty if there is heterogeneity in the loadings on common risk factors. Similarly, cross-sectional dispersion in firm-level profits, sales, and productivity can fluctuate over the business cycle merely because there is heterogeneity in the cyclicalities of firms’ business activity.³

This paper provides new measures of uncertainty and relates them to macroeconomic activity. Our goal is to provide superior econometric estimates of uncertainty that are as free as possible both from the structure of specific theoretical models, and from dependencies on any single (or small number) of observable economic indicators. We start from the premise that what matters for economic decision making is not whether particular economic indicators have become more or less variable or disperse *per se*, but rather whether the economy has become more or less *predictable*; that is, less or more uncertain.

To formalize our notion of uncertainty, let us define h -period ahead uncertainty in the variable $y_{jt} \in Y_t = (y_{1t}, \dots, y_{N_y t})'$, denoted by $\mathcal{U}_{jt}^y(h)$, to be the conditional volatility of the purely unforecastable component of the future value of the series. Specifically,

³Abraham and Katz (1986) also suggested that cross-section variation in employment could vary over the business cycle because of heterogeneity across firms.

$$\mathcal{U}_{jt}^y(h) \equiv \sqrt{E \left[(y_{jt+h} - E[y_{jt+h}|I_t])^2 | I_t \right]} \quad (2.1)$$

where the expectation $E(\cdot|I_t)$ is taken with respect to information I_t available to economic agents at time t .⁴ If the expectation today (conditional on all available information) of the squared error in forecasting y_{jt+h} rises, uncertainty in the variable increases. A measure, or index, of *macroeconomic uncertainty* can then be constructed by aggregating individual uncertainty at each date using aggregation weights w_j :

$$\mathcal{U}_t^y(h) \equiv \text{plim}_{N_y \rightarrow \infty} \sum_{j=1}^{N_y} w_j \mathcal{U}_{jt}^y(h) \equiv E_w[\mathcal{U}_{jt}^y(h)]. \quad (2.2)$$

We use the terms macro and aggregate uncertainty interchangeably.

We emphasize two features of these definitions. First, we distinguish between uncertainty in a series y_{jt} and its conditional volatility. The proper measurement of uncertainty requires removing the forecastable component $E[y_{jt+h}|I_t]$ before computing conditional volatility. Failure to do so will lead to estimates that erroneously categorize forecastable variations as “uncertain.” Thus, uncertainty in a series is *not* the same as the conditional volatility of the raw series where for example a constant mean is removed: it is important to remove the entire forecastable component. While this point may seem fairly straightforward, it is worth noting that almost all measures of stock market volatility (realized or implied) or cross-sectional dispersion currently used in the literature do not take this into account.⁵ We show below that this matters empirically for a large number of series, including the stock market.

⁴A concept that is often related to uncertainty is risk. In a finance context, risk is often measured by conditional covariance of returns with the stochastic discount factor in equilibrium models. This covariance can in turn be driven by conditional volatility in stock returns. [Andersen et al. \(2013\)](#) provide a comprehensive review of the statistical measurement of the conditional variance of financial returns. Uncertainty as defined here is (see discussion below) distinct from conditional volatility but could be one of several reasons why the conditional variances and covariances of returns vary.

⁵Two exceptions are [Gilchrist et al. \(2014\)](#), who use the financial factors developed by [Fama and French \(1992\)](#) to control for common forecastable variation in their measure of realized volatility, and [Bachmann et al. \(2013\)](#), who use subjective forecasts of analysts.

Second, macroeconomic uncertainty is not equal to the uncertainty in any single series y_{jt} . Instead, it is a measure of the common variation in uncertainty across many series. This is important because uncertainty-based theories of the business cycle typically require the existence of common (often countercyclical) variation in uncertainty across large numbers of series. Indeed, in many models of the literature cited above, macroeconomic uncertainty is either directly presumed by introducing stochastic volatility into aggregate shocks (e.g., shocks to aggregate technology, representative-agent preferences, monetary or fiscal policy), or indirectly imposed by way of a presumed countercyclical component in the volatilities of individual firm- or household-level disturbances.⁶ This common variation is critical for the study of business cycles because if the variability of the idiosyncratic shock were entirely idiosyncratic, it would have no influence on macroeconomic variables. If these assumptions are correct, we would expect to find evidence of an aggregate uncertainty factor, or a common component in uncertainty fluctuations that affects many series, sectors, markets, and geographical regions at the same time.

The objective of our paper is therefore to obtain estimates of (2.1) and (2.2). To make these measures of uncertainty operational, we require three key ingredients. First, we require an estimate of the forecast $E[y_{jt+h}|I_t]$. For this, we form factors from a large set of predictors $\{X_{it}\}$, $i = 1, 2, \dots, N$, whose span is as close to I_t as possible. Using these factors, we then approximate $E[y_{jt+h}|I_t]$ by a diffusion index forecast ideal for data-rich environments. An important aspect of this data-rich approach is that the diffusion indices (or common factors) can be treated as known in the subsequent analysis. Second, defining the h -step-ahead forecast error to be $V_{jt+h}^y \equiv y_{jt+h} - E[y_{jt+h}|I_t]$, we require an estimate of the conditional (on time t information) volatility of this error, $E[(V_{t+h}^y)^2|I_t]$. For this, we specify a parametric stochastic volatility model for both the one-step-ahead prediction errors in y_{jt} and the analogous forecast errors for the factors. These volatility estimates are used to recursively compute the

⁶See, e.g., Bloom (2009), Arellano et al. (2012), Bloom et al. (2012), Gilchrist et al. (2014), Schaal (2011), Bachmann and Bayer (2011)). Herskovic et al. (2014) find evidence of a common component in idiosyncratic volatility of firm-level cash-flow growth and returns.

values of $E[(V_{t+h}^y)^2|I_t]$ for $h > 1$. As we show below, this procedure takes into account an important property of multistep-ahead forecasts, namely that time-varying volatility in the errors of the predictor variables creates additional unforecastable variation in y_{jt+h} (above and beyond that created by stochastic volatility in the one-step-ahead prediction error), and contributes to its uncertainty. The third and final ingredient is an estimate of macroeconomic uncertainty $\mathcal{U}_t^y(h)$ constructed from the individual uncertainty measures $\mathcal{U}_{jt}^y(h)$. Our base-case estimate of $\mathcal{U}_t^y(h)$ is the equally-weighted average of individual uncertainties. It is also possible to let the weights be constructed so that macroeconomic uncertainty is interpreted as the common (latent) factor in the individual measures of uncertainty.

We estimate measures of macroeconomic uncertainty from two post-war datasets of economic activity. The first *macro* dataset is monthly and uses the information in hundreds of macroeconomic and financial indicators. The second *firm level* dataset is quarterly and consists of 155 firm-level observations on profit growth normalized by sales. We will refer to estimates of macro uncertainty based on the monthly series as *common macro uncertainty* whereas estimates of macro uncertainty based on the quarterly firm-level dataset will be referred to as *common firm-level uncertainty*.

Our main results may be summarized as follows. We find significant independent variation in our estimates of uncertainty as compared to commonly used proxies for uncertainty. An important finding is that our estimates imply far fewer large uncertainty episodes than what is inferred from all of the commonly used proxies we study. For example, consider the 17 uncertainty dates defined in Bloom (2009) as events associated with stock market volatility in excess of 1.65 standard deviations above its trend. By contrast, in a sample extending from 1960:07 to 2011:12, our measure of macro uncertainty exceeds (or come close to exceeding) 1.65 standard deviations from its mean a total of only 49 (out of 618) months, each of which are bunched into three deep recession episodes discussed below. Moreover, our estimate of macroeconomic uncertainty is far more persistent than stock market volatility: the response of macro uncertainty to its own innovation from an autoregression has a half life of 53 months;

the comparable figure for stock market volatility is 4 months. Qualitatively, these results are similar for our measures of common firm-level uncertainty in profit growth rates. Taken together, the findings imply that most movements in common uncertainty proxies, such as stock market volatility (the most common), and measures of cross-sectional dispersion, are not associated with a broad-based movement in economic uncertainty as defined in (2.2). This is important because it suggests that much of the variation in common uncertainty proxies is not driven by uncertainty.

So how important is time-varying economic uncertainty, and to what extent is it dynamically correlated with macroeconomic fluctuations? Our estimates of macro uncertainty reveal three big episodes of uncertainty in the post-war period: the months surrounding the 1973-74 and 1981-82 recessions and the Great Recession of 2007-09. Averaged across all uncertainty forecast horizons, the 2007-09 recession represents the most striking episode of heightened uncertainty since 1960, with the 1981-82 recession a close second. Large positive innovations to macro uncertainty lead to a sizable and protracted decline in real activity (production, hours, employment). These effects are larger and far more persistent and do not exhibit the “overshooting” pattern found previously when stock market volatility is used to proxy for uncertainty. Using an eleven variable monthly macro vector autoregression (VAR) and a recursive identification procedure with uncertainty placed last, we find that common macro uncertainty shocks account for up to 29% of the forecast error variance in industrial production, depending on the VAR forecast horizon. By contrast, stock market volatility explains at most 7%. To form another basis for comparison, shocks to the federal funds rate (a common proxy for unanticipated shifts in monetary policy) explain (at most) the same amount of forecast error variance in production as does macroeconomic uncertainty, despite uncertainty being placed last in the VAR. Finally, we ask how much each series’ time-varying *individual* uncertainty is explained by time-varying *macro* uncertainty and find that the role of the latter is strongly countercyclical, roughly doubling in importance during recessions.

These results underscore the importance of considering how aggregate uncertainty is

measured when assessing its relationship with the macroeconomy. In particular, our estimates imply that quantitatively important uncertainty episodes occur far more infrequently than what is indicated from common uncertainty proxies, but that when they *do* occur, they display larger and more persistent correlations with real activity. Indeed, the deepest, most protracted recessions in our sample are associated with large increases in estimated uncertainty, while more modest reductions in real activity are not. By contrast, common uncertainty proxies are less persistent and spike far more frequently, often in non-recession periods, or in periods of relative macroeconomic quiescence.

While we find that increases in uncertainty are associated with large declines in real activity, we caution that our results are silent on whether uncertainty is the cause or effect of such declines. Our goal is to develop a defensible measure of time-varying macro uncertainty that can be tracked over time and related to fluctuations in real activity and asset markets. Our estimates do, however, imply that the economy is objectively less predictable in recessions than it is in normal times. This result is not a statement about changing subjective perceptions of uncertainty in recessions as compared to booms. Any theory for which uncertainty is entirely the effect of recessions would need to be consistent with these basic findings.

In this way, our estimates provide a benchmark with which to evaluate theories where uncertainty plays a role in business cycles. Uncertainty as defined in this paper only requires evaluation of the h step ahead conditional expectation and conditional volatility of the variable in question and so can be computed for any number of endogenous variables in a dynamic, stochastic, general equilibrium (DSGE) model. Moreover, these statistics can be computed from within the model regardless of whether the theory implies that uncertainty is the cause or effect of recessions. A comparison of the uncertainty implied by the model and the data can be used to evaluate DSGE models that feature uncertainty.

The rest of this paper is organized as follows. Section [2.2](#) reviews related empirical literature on uncertainty in more detail. Section [2.3](#) outlines the econometric framework employed

in our study, and describes how our measures of uncertainty are constructed. Section 2.4 describes the data and empirical implementation. Section 2.5 presents our common macro uncertainty estimates, compares our measure to other proxies of uncertainty used in the literature, and considers the dynamic relationship between macro uncertainty and variables such as production and employment. Section 2.6 performs a similar analysis for our estimates of common firm-level uncertainty. Section 2.7 summarizes and concludes.

To conserve space, a large amount of supplementary material for this paper appears in Appendix B. This appendix has two parts. The first part provides results from a large number of robustness exercises designed to check the sensitivity of our results to various assumptions (see description below). The second part is a data appendix that contains details on the construction of all data used in this study, including data sources.

2.2 Related Empirical Literature

The literature on measuring uncertainty is still in its infancy. Existing research has primarily relied on measures of volatility and dispersion as proxies of uncertainty. In his seminal work, Bloom (2009) found a strong countercyclical relationship between real activity and uncertainty as proxied by stock market volatility. His VAR estimates suggest that uncertainty has an impact on output and employment in the six months after an innovation in these measures, with a rise in volatility at first depressing real activity and then increasing it, leading to an over-shoot of its long-run level, consistent with the predictions of models with uncertainty as a driving force of macroeconomic fluctuations. Bloom et al. (2012) also documented a relation between real activity and uncertainty as proxied by dispersion in firm-level earnings, industry-level earnings, total factor productivity, and the predictions of forecasters. A recurring feature of these studies is that the uncertainty proxies are strongly countercyclical.

While these analyses are sensible starting places and important cases to understand, we

emphasize here that the measures of dispersion and stock market volatility studied may or may not be tightly linked to true economic uncertainty. Indeed, one of the most popular proxies for uncertainty is closely related to financial market volatility as measured by the VIX, which has a large component that appears driven by factors associated with time-varying risk-aversion rather than economic uncertainty (Bekaert et al., 2013).

A separate strand of the literature focuses on cross-sectional dispersion in N_A analysts' or firms' subjective expectations as a measure of uncertainty:

$$\mathcal{D}_{jt}^A(h) = \sqrt{\sum_{k=1}^{N_{A_t}} w_k^A \left[(y_{jt+h} - E(y_{jt+h}|I_{A_k,t}))^2 | I_{A_k,t} \right]^2}$$

where $I_{A_k,t}$ is the information of agent k at time t , and w_k^A is the weight applied to agent k . One potential advantage of using $\mathcal{D}_{jt}^A(h)$ as a proxy for uncertainty is that it treats the conditional forecast of y_{jt+h} as an observable variable, and therefore does not require estimation of $E[y_{t+h}|I_{A_k,t}]$. Bachmann and Bayer (2011) follow this approach using a survey of German firms and argue that uncertainty appears to be more an outcome of recessions than a cause, contrary to the predictions of theoretical models such as Bloom (2009) and Bloom et al. (2012). D'Amico and Orphanides (2008) is an earlier project that studies various measures of analyst uncertainty and disagreement from the Survey of Professional Forecasters. While analysts' forecasts are interesting in their own right, there are several known drawbacks in using them to measure uncertainty. First, subjective expectations are only available for a limited number of series. For example, of the 132 monthly macroeconomic series we will consider in this paper, not even one-fifth have corresponding expectations series. Second, it is not clear that the responses elicited from these surveys accurately capture the conditional expectations of the economy as a whole. The respondents typically sampled are practitioner forecasters; some analysts' forecasts are known to display systematic biases and omit relevant forecasting information (So, 2013), and analysts may have pecuniary incentives to bias their forecasts in a way that economic agents would not. Third, disagreement in

survey forecasts could be more reflective of differences in opinion than of uncertainty (e.g. [Diether et al., 2002](#); [Mankiw et al., 2003](#)). As discussed above, it could also reflect differences in firm’s loadings on aggregate shocks in the absence of aggregate or idiosyncratic time-varying volatility. Fourth, [Lahiri and Sheng \(2010\)](#) show that, even if forecasts are unbiased, disagreement in analysts’ point forecasts does not equal (average across analysts) forecast error uncertainty unless the variance of accumulated aggregate shocks over the forecast horizon is zero. They show empirically using the Survey of Professional Forecasters that the variance of the accumulated aggregate shocks can drive a large wedge between uncertainty and disagreement in times of important economic change, or whenever the forecast horizon is not extremely short. [Bachmann and Bayer \(2011\)](#) acknowledge these problems and are careful to address them by using additional proxies for uncertainty, such as an ex-post measure of forecast error variance based on the survey expectations. A similar approach is taken in [Scotti \(2013\)](#) who studies series for which real-time data are available. Whereas these studies focus on variation in outcomes around subjective survey expectations of relatively few variables, we focus on uncertainty around objective statistical forecasts for hundreds of economic series.

Our uncertainty measure is also different from proxies based on the unconditional cross-section dispersion of a particular variable:

$$\mathcal{D}_{jt}^B = \sqrt{\frac{1}{N_B} \sum_{k=1}^{N_B} \left[\left(y_{jkt} - \frac{1}{N_B} \sum_i y_{jit} \right)^2 \right]} \quad (2.3)$$

where y_{jkt} is a variable indexed by j (e.g., firm-level profits studied in [Bloom, 2009](#)) for firm k , and N_B is the sample size of firms reporting profits. Notably, this dispersion has no forward looking component; it is the same for all horizons. This measure suffers from the same drawback as $\mathcal{D}_{jt}^A(h)$, namely that it can fluctuate without any change in uncertainty if there is heterogeneity in the cyclicalities of firms business activity.

[Carriero et al. \(2012\)](#) consider common sources of variation in the residual volatilities

of a Bayesian Vector Autoregression (VAR). This investigation differs from ours in several ways: their focus is on small-order VARs (e.g., 4 or 8 variables) and residual volatility, which corresponds to our definition of uncertainty only when $h = 1$; our interest is in measuring the prevalence of uncertainty across the entire macroeconomy. Their estimation procedure presumes that individual volatilities only have common shocks, and it is not possible for some series to have homoskedastic shocks while others have heteroskedastic ones. We find a large idiosyncratic component in individual volatilities, the magnitude of which varies across series.

An important unresolved issue for empirical analysis of uncertainty concerns the persistence of uncertainty shocks. In models studied by [Bloom et al. \(2012\)](#), for example, recessions are caused by an increase in uncertainty, which in turn causes a drop in productivity growth. But other researchers who have studied models where uncertainty plays a key role (e.g. [Schaal, 2011](#)) have argued that empirical proxies for uncertainty, such as the cross-sectional dispersion in firms' sales growth, are not persistent enough to explain the prolonged levels of unemployment that have occurred during and after some recessions, notably the 2007-2009 recession and its aftermath. Here we provide new measures of uncertainty and its persistence, finding that they are considerably more persistent than popular proxies such as stock market volatility and measures of dispersion.

2.3 Econometric Framework

We now turn to a description of our econometric framework. A crucial first step in our analysis is to replace the conditional expectation in (2.1) by a forecast, from which we construct the forecast error that forms the basis of our uncertainty measures. In order to identify a true forecast error, it is important that our predictive model be as rich as possible, so that our measured forecast error is purged of predictive content. A standard approach is to select a set of K predetermined conditioning variables given by the $K \times 1$ vector \mathbf{W}_t , and

then estimate

$$y_{t+1} = \beta' \mathbf{W}_t + \epsilon_{t+1} \quad (2.4)$$

by least squares. The one period forecast is $\hat{y}_{t+1|t} = \hat{\beta}' \mathbf{W}_t$ where $\hat{\beta}$ is the least squares estimate of β .⁷ An omitted-information bias may arise if economic agents such as financial market participants have more information than that in the conditioning variables. Indeed, recent work finds that forecasts of both real activity and financial returns are substantially improved by augmenting best-fitting conventional forecasting equations with common factors estimated from large datasets.⁸

To address this problem, we use the method of diffusion index forecasting whereby a relatively small number of factors estimated from a large number of economic time series are augmented to an otherwise standard forecasting model. The omitted information problem is remedied by including estimated factors, and possibly non-linear functions of these factors or factors formed from non-linear transformations of the raw data, in the forecasting model. This eliminates the arbitrary reliance on a small number of exogenous predictors and enables the use of information in a vast set of economic variables that are more likely to span the unobservable information sets of economic agents. Diffusion index forecasts are increasingly used in data rich environments. Thus we only generically highlight the forecasting step and focus instead on construction of uncertainty, leaving details about estimation of the factors to Appendix B.

⁷ But we often have substantially more information available for the prediction exercise than what is contained in a small number of predictor variables. Suppose we observe a $T \times N$ panel of data with elements $X_{it}, i = 1, \dots, N, t = 1, \dots, T$, where the cross-sectional dimension, N , is large, and possibly larger than the number of time periods, T . The computational challenge is that there are potentially 2^N possible combinations of variables to consider, and estimation is not even feasible when $n + K > T$ for any subset of (X_{1t}, \dots, X_{NT}) of size n . In practice, researchers are forced to choose among a few conditioning variables to overcome these problems.

⁸See, for example, [Stock and Watson \(2002b, 2006\)](#), and [Ludvigson and Ng \(2007, 2009\)](#). This problem is especially important in our exercise since relevant information not used to form forecasts will lead to spurious estimates of uncertainty and its dynamics.

2.3.1 Construction of Forecast Uncertainty

Let $\mathbf{X}_t = (X_{1t}, \dots, X_{Nt})'$ generically denote the predictors available for analysis. It is assumed that \mathbf{X}_t has been suitably transformed (such as by taking logs and differencing) so as to render the series stationary. We assume that X_{it} has an approximate factor structure taking the form

$$X_{it} = \Lambda_i^{F'} \mathbf{F}_t + e_{it}^X, \quad (2.5)$$

where \mathbf{F}_t is an $r_F \times 1$ vector of latent common factors, Λ_i^F is a corresponding $r_F \times 1$ vector of latent factor loadings, and e_{it}^X is a vector of idiosyncratic errors. In an *approximate* dynamic factor structure, the idiosyncratic errors e_{it}^X are permitted to have a limited amount of cross-sectional correlation. Importantly, the number of factors r_F is significantly smaller than the number of series, N .

Let y_{jt} generically denote a series that we wish to compute uncertainty in and whose value in period $h \geq 1$ is estimated from a factor augmented forecasting model

$$y_{jt+1} = \phi_j^y(L)y_{jt} + \gamma_j^F(L)\hat{\mathbf{F}}_t + \gamma_j^W(L)\mathbf{W}_t + v_{jt+1}^y \quad (2.6)$$

where $\phi_j^y(L)$, $\gamma_j^F(L)$, and $\gamma_j^W(L)$ are finite-order polynomials in the lag operator L of orders p_y , p_F , and p_W , respectively, the elements of the vector $\hat{\mathbf{F}}_t$ are consistent estimates of a rotation of \mathbf{F}_t , and the r_w dimensional vector \mathbf{W}_t contains additional predictors that will be specified below. An important feature of our analysis is that the one-step-ahead prediction error of y_{jt+1} , and of each factor $F_{k,t+1}$ and additional predictor $W_{h,t+1}$, is permitted to have time-varying volatility σ_{jt+1}^y , σ_{kt+1}^F , σ_{ht+1}^W , respectively. This feature generates time-varying uncertainty in the series y_{jt} .⁹

When the factors have autoregressive dynamics, a more compact representation of the system above is the *factor augmented vector autoregression* (FAVAR). Let $\mathbf{Z}_t \equiv (\hat{\mathbf{F}}_t', \mathbf{W}_t')'$ be a $r = r_F + r_W$ vector which collects the r_F estimated factors and r_W additional predictors,

⁹Define $z_{jt-1} \equiv \gamma_j^F(L)\hat{\mathbf{F}}_{t-1} + \gamma_j^W(L)\mathbf{W}_{t-1}$, and $q = \max(p_y, p_F, p_W, h)$. Then (2.6) can be written in

and define $\mathcal{Z}_t \equiv (\mathbf{Z}'_t, \dots, \mathbf{Z}'_{t-q+1})'$. Also let $Y_{jt} = (y_{jt}, y_{jt-1}, \dots, y_{jt-q+1})'$. Then forecasts for any $h > 1$ can be obtained from the FAVAR system, stacked in first-order companion form:

$$\begin{pmatrix} \mathcal{Z}_t \\ Y_{jt} \end{pmatrix}_{(r+1)q \times 1} = \begin{pmatrix} \underbrace{\Phi^{\mathcal{Z}}}_{qr \times qr} & \underbrace{\mathbf{0}}_{qr \times q} \\ \underbrace{\Lambda'_j}_{q \times qr} & \underbrace{\Phi^Y_j}_{q \times q} \end{pmatrix} \begin{pmatrix} \mathcal{Z}_{t-1} \\ Y_{jt-1} \end{pmatrix} + \begin{pmatrix} \mathcal{V}^{\mathcal{Z}}_t \\ \mathcal{V}^Y_{jt} \end{pmatrix} \quad (2.8)$$

$$\mathcal{Y}_{jt} = \Phi^{\mathcal{Y}}_j \mathcal{Y}_{jt-1} + \mathcal{V}^{\mathcal{Y}}_{jt},$$

where Λ'_j and Φ^Y_j are functions of the coefficients in the lag polynomials in (2.6), $\Phi^{\mathcal{Z}}$ stacks the autoregressive coefficients of the components of \mathcal{Z}_t .¹⁰ By the assumption of stationarity, the largest eigenvalue of $\Phi^{\mathcal{Y}}_j$ is less than one and, under quadratic loss, the optimal h -period forecast is the conditional mean:

$$E_t \mathcal{Y}_{jt+h} = (\Phi^{\mathcal{Y}}_j)^h \mathcal{Y}_{jt}.$$

first-order companion form as

$$\begin{pmatrix} y_{jt} \\ y_{jt-1} \\ y_{jt-2} \\ \vdots \\ y_{jt-q+1} \end{pmatrix} = \begin{pmatrix} \phi_{j1}^y & \phi_{j2}^y & \phi_{j3}^y & \cdot & \cdot & \phi_{jp_y}^y & \mathbf{0}_{1 \times (q-p_y)} \\ 1 & 0 & 0 & \cdot & \cdot & \cdot & \cdot \\ 0 & 1 & 0 & \cdot & \cdot & \cdot & \cdot \\ \cdot & & \cdot & & & & \cdot \\ \cdot & & & \cdot & & & \cdot \\ \cdot & & & & \cdot & & \cdot \\ 0 & 0 & 0 & \cdot & \cdot & 1 & \mathbf{0}_{1 \times (q-p_y)} \end{pmatrix} \begin{pmatrix} y_{jt-1} \\ y_{jt-2} \\ y_{jt-3} \\ \vdots \\ y_{jt-q} \end{pmatrix} + \begin{pmatrix} z_{jt-1} \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + \begin{pmatrix} v_{jt}^y \\ 0 \\ 0 \\ i \\ \cdot \\ \cdot \\ 0 \end{pmatrix} \quad (2.7)$$

$$Y_{jt} = \Phi^Y_j Y_{t-1} + Z_{jt-1} + \mathcal{V}^Y_{jt},$$

where $\mathbf{0}_{1 \times (q-p_y)}$ is an additional vector of zeros if $p_y < q$. The vector denoted \mathcal{V}^Y_{jt} above should be distinguished from V^Y_{jt+h} , the h period ahead forecast error for the scalar variable y_{jt} that will be explored below.

¹⁰The above specification assumes that the coefficients are time-invariant. Cogley and Sargent (2005) among others have found important variation in VAR coefficients. Dynamic factor models are somewhat more robustness against temporal parameter instability than small forecasting models (Stock and Watson, 2002b). The reason is that such instabilities can “average out” in the construction of common factors if the instability is sufficiently dissimilar from one series to the next. Nonetheless, as a robustness check, an uncertainty measure is also constructed using recursive out-of-sample forecasts errors and will be discussed below.

The forecast error variance at t is

$$\Omega_{jt}^{\mathcal{Y}}(h) \equiv E_t [(\mathcal{Y}_{jt+h} - E_t \mathcal{Y}_{jt+h})(\mathcal{Y}_{jt+h} - E_t \mathcal{Y}_{jt+h})'] .$$

Time variation in the mean squared forecast error in general arises from the fact that shocks to *both* y_{jt} and the predictors \mathbf{Z}_t may have time-varying variances. We now turn to these implications. Note first that when $h = 1$,

$$\Omega_{jt}^{\mathcal{Y}}(1) = E_t(\mathcal{V}_{jt+1}^{\mathcal{Y}} \mathcal{V}_{jt+1}^{\mathcal{Y}'}) . \quad (2.9)$$

For $h > 1$, the forecast error variance of \mathcal{Y}_{jt+h} evolves according to

$$\Omega_{jt}^{\mathcal{Y}}(h) = \Phi_j^{\mathcal{Y}} \Omega_{jt}^{\mathcal{Y}}(h-1) \Phi_j^{\mathcal{Y}'} + E_t(\mathcal{V}_{jt+h}^{\mathcal{Y}} \mathcal{V}_{jt+h}^{\mathcal{Y}'}) . \quad (2.10)$$

As $h \rightarrow \infty$ the forecast is the unconditional mean and the forecast error variance is the unconditional variance of \mathcal{Y}_{jt} . This implies that $\Omega_{jt}^{\mathcal{Y}}(h)$ is less variable as h increases.

We are interested in the expected forecast uncertainty of the scalar series y_{jt+h} given information at time t , denoted $\mathcal{U}_{jt}^y(h)$. This is the square-root of the appropriate entry of the forecast error variance $\Omega_{jt}^{\mathcal{Y}}(h)$. With $\mathbf{1}_j$ being a selection vector,

$$\mathcal{U}_{jt}^y(h) = \sqrt{\mathbf{1}_j' \Omega_{jt}^{\mathcal{Y}}(h) \mathbf{1}_j} . \quad (2.11)$$

To estimate macro (economy-wide) uncertainty, we form weighted averages of individual uncertainty estimates:

$$\sum_{j=1}^{N_y} w_j \mathcal{U}_{jt}^y(h) .$$

A simple weighting scheme is to give every series the equal weight of $w_j = 1/N_y$. If individual uncertainty has a factor structure, the weights can be defined by the eigenvector corresponding to the largest eigenvalue of the $N_y \times N_y$ covariance matrix of the matrix of

individual uncertainty. We discuss our weighting schemes for measuring macro uncertainty further below.

2.3.2 Time-varying Uncertainty: A Statistical Decomposition

In this subsection we show how stochastic volatility in the predictors Z and in y_j contribute to its h period ahead uncertainty. The choice of stochastic volatility is important because it permits the construction of a shock to the second moment that is independent of innovations to y_j itself. This is consistent with much of the theoretical literature on uncertainty which presumes the existence of a uncertainty shock that independently affects real activity. GARCH type models (for example) do not share this feature and instead have a shock that is *not* independent from innovations to y_j .

Consider first the factors \mathbf{F}_t (the argument for \mathbf{W}_t is similar). Suppose that each element of \mathbf{F}_t is serially correlated and well represented by a univariate AR(1) model (dropping the subscript that indexes the factor in question for simplicity):

$$\mathbf{F}_t = \Phi^F \mathbf{F}_{t-1} + \mathbf{v}_t^F.$$

If v_t^F was a martingale difference with constant variance $(\sigma^F)^2$, the forecast error variance $\Omega^F(h) = \Omega^F(h-1) + (\Phi^F)^{2(h-1)}(\sigma^F)^2$ increases with h but is the same for all t . We allow the shocks to \mathbf{F} to exhibit time-varying stochastic volatility, i.e. $v_t^F = \sigma_t^F \varepsilon_t^F$ where log volatility has an autoregressive structure:

$$\log(\sigma_t^F)^2 = \alpha^F + \beta^F \log(\sigma_{t-1}^F)^2 + \tau^F \eta_t^F, \quad \eta_t^F \stackrel{Fiid}{\sim} N(0, 1).$$

The stochastic volatility model allows for a shock to the second moment that is independent

of the first moment, consistent with theoretical models of uncertainty. The model implies

$$E_t(\sigma_{t+h}^F)^2 = \exp \left[\alpha^F \sum_{s=0}^{h-1} (\beta^F)^s + \frac{(\tau^F)^2}{2} \sum_{s=0}^{h-1} (\beta^F)^{2(s)} + (\beta^F)^h \log(\sigma_t^F)^2 \right].$$

Since $\epsilon_t^F \stackrel{iid}{\sim} (0, 1)$ by assumption, $E_t(v_{t+h}^F)^2 = E_t(\sigma_{t+h}^F)^2$. This allows us to compute the $h > 1$ forecast error variance for F using the recursion

$$\Omega_t^F(h) = (\Phi^F) \Omega_t^F(h-1) \Phi^{F'} + E_t(v_{t+h}^F v_{t+h}^{F'})$$

with $\Omega_t^F(1) = E_t(v_{t+1}^F)^2$. The h period ahead *predictor uncertainty* at time t is the square root of the h -step forecast error variance of the predictor:

$$\mathcal{U}_t^F(h) = \sqrt{1_F' \Omega_t^F(h) 1_F}$$

where 1_F is an appropriate selection vector. It follows from the determinants of $E(\sigma_{t+h}^F)^2$ that h -period-ahead uncertainty of F_t has a level-effect attributable to α^F (the homoskedastic variation in v_{Ft}), a scale effect attributable to τ^F , with persistence determined by β^F .

To understand how uncertainty in the predictors affect uncertainty in the variable of interest y_j , suppose that the forecasting model for \mathbf{y}_j only has a single predictor $\hat{\mathbf{F}}$ and is given by:

$$y_{jt+1} = \phi_j^y y_{jt} + \gamma_j^F \hat{\mathbf{F}}_t + v_{jt+1}^y$$

where $v_{jt+1}^y = \sigma_{jt+1}^y \varepsilon_{jt+1}^y$ with $\varepsilon_{jt+1}^y \stackrel{iid}{\sim} N(0, 1)$ and

$$\log(\sigma_{jt+1}^y)^2 = \alpha_j^y + \beta_j^y \log(\sigma_{jt}^y)^2 + \tau_j^y \eta_{jt+1}, \quad \eta_{jt+1} \stackrel{iid}{\sim} N(0, 1).$$

When $h = 1$, V_{jt+1}^y coincides with the innovation v_{jt+1}^y which is uncorrelated with the one-step-ahead error in forecasting F_{t+1} , given by $V_{t+1}^F = v_{t+1}^F$. When $h = 2$, the forecast

error for the factor is $V_{t+2}^F = \Phi^F V_{t+1}^F + v_{t+2}^F$. The corresponding forecast error for y_{jt} is:

$$V_{jt+2}^y = v_{jt+2}^y + \phi_j^y V_{jt+1}^y + \gamma_j^F V_{t+1}^F$$

which evidently depends on the one-step-ahead forecasting errors made at time t , but V_{t+1}^y and V_{t+1}^F are uncorrelated. When $h = 3$, the forecast error is

$$V_{jt+3}^y = v_{jt+3}^y + \phi_j^y V_{jt+2}^y + \gamma_j^F V_{t+2}^F$$

which evidently depends on V_{jt+2}^y and V_{t+2}^F . But unlike the $h = 2$ case, the two components V_{jt+2}^y and V_{t+2}^F are now correlated because both depend on V_{t+1}^F .

Therefore, returning to the general case when the predictors are $\mathbf{Z}_t = (\mathbf{F}_t', \mathbf{W}_t')'$ and its lags, h -step-ahead forecast error variance for Y_{jt+h} admits the decomposition:

$$\Omega_{jt}^Y(h) = \underbrace{\Phi_j^Y \Omega_{jt}^Y(h-1) \Phi_j^{Y'}}_{\text{autoregressive}} + \underbrace{\Omega_{jt}^Z(h-1)}_{\text{Predictor}} + \underbrace{E_t(\mathcal{V}_{jt+h}^Y \mathcal{V}_{jt+h}^{Y'})}_{\text{stochastic volatility } Y} + \underbrace{2\Phi_j^Y \Omega_{jt}^{YZ}(h-1)}_{\text{covariance}} \quad (2.12)$$

where $\Omega_{jt}^{YZ}(h) = \text{cov}_t(\mathcal{V}_{jt+h}^Y, \mathcal{V}_{jt+h}^Z)$. The terms in $E(\mathcal{V}_{j,t+h}^Y \mathcal{V}_{j,t+h}^{Y'})$ are computed using the fact that $E_t(v_{jt+h}^y)^2 = E_t(\sigma_{jt+h}^y)^2$, $E_t(v_{t+h}^F)^2 = E_t(\sigma_{t+h}^F)^2$ and $E_t(v_{t+h}^W)^2 = E_t(\sigma_{t+h}^W)^2$.

Time variation in uncertainty can thus be mathematically decomposed into four sources: an autoregressive component, a common factor (predictor) component, a stochastic volatility component, and a covariance term. Representation (2.12), which is equivalent to (2.10) for the subvector \mathbf{Y}_t , makes clear that predictor uncertainty plays an important role via the second term $\Omega_{jt}^Z(h-1)$. It is time-varying because of stochastic volatility in the innovations to the factors and is in general non-zero for multi-step-ahead forecasts, i.e., $h > 1$. The role of stochastic volatility in the series y_j comes through the third term, with the role of the covariance between the forecast errors of the series and the predictors coming through the last term. Computing the left-hand-side therefore requires estimates of stochastic volatility in the residuals of every series \mathbf{y}_j , and in every component of the predictor variables \mathbf{Z} .

2.4 Empirical Implementation and Macro Data

Our empirical analysis forms forecasts and common uncertainty from two datasets. The first dataset, denoted \mathbf{X}^m , is an updated version of the 132 mostly macroeconomic series used in [Ludvigson and Ng \(2010\)](#). The 132 macro series in \mathbf{X}^m are selected to represent broad categories of macroeconomic time series: real output and income, employment and hours, real retail, manufacturing and trade sales, consumer spending, housing starts, inventories and inventory sales ratios, orders and unfilled orders, compensation and labor costs, capacity utilization measures, price indexes, bond and stock market indexes, and foreign exchange measures. The second dataset, denoted \mathbf{X}^f , is an updated monthly version of the of 147 financial time series used in [Ludvigson and Ng \(2007\)](#). The data include valuation ratios such as the dividend-price ratio and earnings-price ratio, growth rates of aggregate dividends and prices, default and term spreads, yields on corporate bonds of different ratings grades, yields on Treasuries and yield spreads, and a broad cross-section of industry, size, book-market, and momentum portfolio equity returns. A detailed description of the series is given in the [Appendix B](#). Both of these datasets span the period 1960:01-2011:12. After lags in the FAVAR and transformations of the raw data, we construct uncertainty estimates for the period 1960:07-2011:12, or 618 observations.

We combine the macro and financial monthly datasets together into one large “macroeconomic dataset” (\mathbf{X}) to estimate forecasting factors in these 132+147=279 series. However, we estimate macroeconomic uncertainty $\mathcal{U}_t^y(h)$ from the individual uncertainties in the 132 macro series only. Uncertainties in the 147 financial series are not computed because \mathbf{X}^m already includes a number of financial indicators. To obtain a broad-based measure of uncertainty, it is desirable not to over-represent the financial series, which are far more volatile than the macro series and can easily dominate the aggregate uncertainty index.¹¹

The stochastic volatility parameters $\alpha_j, \beta_j, \tau_j$ are estimated from the least square residuals

¹¹The macro dataset already contains some 25 financial indicators. If we include the additional 147 indicators in our uncertainty index, their greater volatility will dominate the uncertainty measure and we will get back a aggregate financial market volatility variable as uncertainty.

of the forecasting models using Markov chain Monte Carlo (MCMC) methods.¹² In the base-case, the average of these model parameters over the MCMC draws are used to estimate $\mathcal{U}_{jt}^y(h)$. Simple averaging is used to obtain an estimate of h period macro uncertainty denoted

$$\bar{\mathcal{U}}_t^y(h) = \frac{1}{N^y} \sum_{j=1}^{N^y} \hat{\mathcal{U}}_{jt}^y(h), \quad (2.13)$$

where the “hat” indicates the estimated value of \mathcal{U}_{jt}^y . This measure of average uncertainty does not impose any structure on the individual uncertainties above and beyond the assumed assumptions on the latent volatility process.

As an alternative to equally weighting the individual uncertainty estimates, we also construct a latent common factor estimate of macro uncertainty as the first principal component of the covariance matrix of individual uncertainties, denoted $\bar{\mathcal{U}}_t(h)$. To ensure that the latent uncertainty factor is positive, the method of principal components is applied to the logarithm of the individual uncertainty estimates and then rescaled. Its construction is detailed in the Appendix B.

Throughout, the factors in the forecasting equation are estimated by the method of static principal components (PCA). [Bai and Ng \(2006\)](#) show that if $\sqrt{T}/N \rightarrow 0$, the estimates \hat{F}_t can be treated as though they were observed in the subsequent forecasting regression. The defining feature of a model with r_F factors is that the r_F largest population eigenvalues should increase as N increases, while the $N - r_F$ eigenvalues should be bounded. The criterion of [Bai and Ng \(2002\)](#) suggests $r_F = 12$ forecasting factors \mathbf{F}_t for the combined datasets \mathbf{X}^m and \mathbf{X}^f explaining about 54% of the variation in the 279 series, with the first three factors accounting for 37%, 8%, 3%, respectively. The first factor loads heavily on stock market portfolio returns (such as size and book-market portfolio returns), the excess stock market return, and the log dividend-price ratio. The second factor loads heavily on measures of real

¹²We use the STOCHVOL package in R, which implements the ancillarity-sufficiency interweaving strategy as discussed in [Kastner and Frühwirth-Schnatter \(2014\)](#) which is less sensitive to whether the mean of the volatility process is in the observation or the state equation. Earlier versions of this paper implements the algorithm of [Kim et al. \(1998\)](#) using our own MATLAB code.

activity, such as manufacturing production, employment, total production and employment, and capacity utilization. The third factor loads heavily on risk and term spreads in the bond market.

The potential predictors in the forecasting model are $\hat{\mathbf{F}}_t = (\hat{F}_{1t}, \dots, \hat{F}_{r_F t})'$ and \mathbf{W}_t , where \mathbf{W}_t consists of squares of the first component of $\hat{\mathbf{F}}_t$, and factors in X_{it}^2 collected into the $N_G \times 1$ vector $\hat{\mathbf{G}}_t$. These quadratic terms in \mathbf{W}_t are used to capture possible non-linearities and any effect that conditional volatility might have on the conditional mean function. Following [Bai and Ng \(2008\)](#), the predictors ultimately used are selected so as to insure that only those likely to have significant incremental predictive power are included. To do so, we apply a hard thresholding rule using a conservative t test to retain those \mathbf{F}_t and \mathbf{W}_t that are statistically significant.¹³ The most frequently selected predictors are \hat{F}_{2t} , a “real” factor highly correlated with measures of industrial production and employment, \hat{F}_{12t-1} , highly correlated with lagged hours, \hat{F}_{4t} , highly correlated with measures of inflation, and \hat{F}_{10t} , highly correlated with exchange rates. Four lags of the dependent variable are always included in the predictive regressions.

Before describing the results, we comment briefly on the question of whether it is desirable for our objective to use so-called “real-time” data, which would restrict the forecasting information set to observations on X_{it} that coincide with the estimated value for this series available at time t from data collection agencies. Such a dataset differs from the final “historical” data on X_{it} because initial estimates of a series are available only with a (typically one month) delay, and earlier available estimates of many series are revised in subsequent months as better estimates become available. In this paper we use the final revised, or historical, data in our estimation, for two reasons. The first is a practical one: our approach calls for a summary statistic of forecasts and therefore uncertainty across many series, requiring far

¹³Specifically, we begin with a set of candidate predictors that includes all the estimated factors in X_{it} (the $\hat{\mathbf{F}}_t$), the first estimated factor in X_{it}^2 (\hat{G}_{1t}), and the square of the first factor in X_{it} (\hat{F}_{1t}^2). We then chose subsets from these by running a regression of y_{it+1} on a constant, four lags of the dependent variable, $\hat{\mathbf{F}}_t$, \hat{F}_{1t}^2 , and \hat{G}_{1t} (no lags). Regressors are retained if they have a marginal t statistic greater than 2.575 in the multivariate forecasting regression of y_{it+1} on the candidate predictors known at time t .

more series than what is in practice available on a real-time data basis.

Second, and more fundamentally, we are interested in forming the most historically accurate estimates of uncertainty at any given point in time in our sample. Restricting information to real-time data is not ideal for this objective because it is likely to be overly restrictive, underestimating the amount of information agents actually had at the time of the forecast. Economic modeling is replete with examples of why this could be so. In representative-agent models, agents typically observe the current aggregate economic state as it occurs. In practice, individuals know their own consumption, incomes, the prices they pay for consumption goods, and probably a good deal about the output of the firm and industries they work in, long before data collection agencies report on these. Even forecasting practitioners can predict a large fraction of a future data release based on current information. In this sense, except for data from asset markets, many of what is called real-time data is not really real-time news, but instead represents newly released information on events that had occurred. Even in heterogeneous-agent models where individuals directly observe only their own economic state variables, the aggregate state upon which their optimization problems depend can typically be well summarized by a few financial market returns that are observable on a timely basis. Partly for this reason, our forecasting equations always include a large number of financial indicators as conditioning variables. The 147 financial data series include many empirical risk-factors for stocks and bonds that we expect to be immediately responsive to any genuine news contained in data releases. These financial indicators can also be expected to respond in real time to disaster-like events (wars, political shocks, natural disasters) that invariably increase uncertainty.¹⁴

¹⁴Baker and Bloom (2013) use disaster-like events as instruments for stock market volatility with the objective of sorting out the causal relationship between uncertainty and economic growth.

2.5 Estimates of Macro Uncertainty

We present estimates of macro uncertainty for three horizons: $h = 1, 3$, and 12 months. Figure 2.1 plots $\bar{U}_t^y(h)$ over time for $h = 1, 3$, and 12, along with the NBER recession dates. The matching horizontal bars correspond to 1.65 standard deviations above the mean for each series. Figure 2.1 shows that macro uncertainty is clearly countercyclical: the correlation of $\bar{U}_t^y(h)$ with industrial production growth is -0.62, -0.61, and -0.57 for $h = 1, 3$, and 12, respectively. While the level of uncertainty increases with h (on average), the variability of uncertainty decreases because the forecast tends to the unconditional mean as the forecast horizon tends to infinity. Macro uncertainty exhibits spikes around the 1973-74 and 1981-82 recessions, as well as the Great Recession of 2007-09.

Looking across all uncertainty forecast horizons $h = 1, 3$, and 12, the 2007-09 recession clearly represents the most striking episode of heightened uncertainty since 1960. The 1981-82 recession is a close second, especially for forecast horizons $h = 3$ and 12. Indeed, for these horizons, these are the only two episodes for which macro uncertainty exceeds 1.65 standard deviations above its mean in our sample. Inclusive of $h = 1$, the three episodes are the only instances in which $\bar{U}_t^y(h)$ exceeds, or comes close to exceeding, 1.65 standard deviation above its mean, implying far fewer uncertainty episodes than other popular proxies for uncertainty, as we show below. Heightened uncertainty is broad-based during these three episodes as the fraction of series with $\hat{U}_{jt}^y(h)$ exceeding their own standard deviation over the full sample are 0.42, 0.61, and 0.51 for 1, 3, and 12 respectively. Further investigation reveals that the three series with the highest uncertainty between 1973:11 and 1975:03 are a producer price index for intermediate materials, a commodity spot price index, and employment in mining. For the 1980:01 and 1982:11 episode, uncertainty is highest for the Fed funds rate, employment in mining, and the 3 months commercial paper rate. Between 2007:12 and 2009:06, uncertainty is highest for the monetary base, non-borrowed reserves and total reserves. These findings are consistent with the historical account of an energy crisis around 1974, a recession of monetary policy origin around 1981, and a financial crisis around 2008 that created challenges for the

operation of monetary policy.




Figure 2.1: Aggregate Uncertainty: $\bar{U}_t^y(h)$ for $h = 1, 3, 12$. Horizontal lines indicate 1.65 standard deviations above the mean of each series. Industrial Production (IP) growth is computed as the 12-month moving average of monthly growth rates (in percent). The data are monthly and span the period 1960:07-2011:12.

Table 2.1 reports summary statistics of $\bar{U}_t^y(1)$.¹⁵ The table reports the first-order autocorrelation coefficient, estimates of the half-life of an aggregate uncertainty innovation from a univariate autoregression (AR) for $\bar{U}_t^y(1)$, estimates of skewness, and kurtosis, and $\text{IP-Corr}(k) = |\text{corr}(\bar{U}_t^y(1), \Delta IP_{t+k})|$ is the (absolute) cross-correlation of $\bar{U}_t^y(1)$ with industrial production growth at different leads and lags, k . Also reported is the maximum of $\text{IP-Corr}(k)$ over k . The same statistics are reported for other uncertainty proxies, discussed below. Several statistical facts about the estimate of aggregate uncertainty $\bar{U}_t^y(1)$ stand out in Table 2.1.

¹⁵The statistics for $\bar{U}_{jt}^y(3)$ and $\bar{U}_{jt}^y(12)$ (not reported) are very similar.

Statistic (Monthly)	VXO	$\mathcal{D}(\text{Returns})$	$\bar{\mathcal{U}}_t^y(1)$
AR(1), Half life	0.85, 4.13	0.70, 1.92	0.99, 53.58
Skewness, Kurtosis	2.18, 11.05	1.30, 5.51	1.81, 7.06
IP-Corr(0)	-0.32	-0.45	-0.62
IP-Corr(12), IP-Corr(-12)	-0.29,-0.11	-0.25,-0.07	-0.44,-0.14
$\max_{k>0}\text{IP-Corr}(k)$	-0.43	-0.47	-0.67
At lag $k =$	6	2	3
$\max_{k<0}\text{IP-Corr}(k)$	-0.30	-0.43	-0.59
At lag $k =$	-1	-1	-1
Statistic (Quarterly)		$\mathcal{D}(\text{Profits})$	$\bar{\mathcal{U}}_t^y(1)$
AR(1), Half life		0.76, 2.55	0.93, 10.23
Skewness, Kurtosis		0.36, 2.47	1.77, 6.75
IP-Corr(0)		-0.37	-0.64
IP-Corr(4), IP-Corr(-4)		-0.11,-0.30	-0.46,-0.16
$\max_{k>0}\text{IP-Corr}(k)$		-0.34	-0.70
At lag $k =$		1	1
$\max_{k<0}\text{IP-Corr}(k)$		-0.37	-0.53
At lag $k =$		-1	-1
Statistic (Semi-Annual)		$\mathcal{D}(\text{Forecasts})$	$\bar{\mathcal{U}}_t^y(1)$
AR(1), Half Life		0.45, 0.86	0.85, 4.18
Skewness, Kurtosis		0.24, 2.19	1.74, 6.41
IP-Corr(0)		-0.41	-0.64
IP-Corr(2), IP-Corr(-2)		-0.16,-0.24	-0.48,-0.16
$\max_{k>0}\text{IP-Corr}(k)$		-0.34	-0.68
At lag $k =$		1	1
$\max_{k<0}\text{IP-Corr}(k)$		-0.35	-0.40
At lag $k =$		-1	-1
Statistic (Annual)		$\mathcal{D}(\text{TFP})$	$\bar{\mathcal{U}}_t^y(1)$
AR(1), Half Life		0.33, 0.63	0.61, 1.40
Skewness, Kurtosis		1.71, 8.56	1.76, 6.24
IP-Corr(0)		-0.55	-0.69
IP-Corr(1), IP-Corr(-1)		-0.33,-0.09	-0.48,-0.16
$\max_{k>0}\text{IP-Corr}(k)$		-0.33	-0.47
At lag $k =$		1	1
$\max_{k<0}\text{IP-Corr}(k)$		-0.24	-0.24
At lag $k =$		-26	-26

Table 2.1: Summary statistics. IP-Corr(k) is the absolute cross-correlation between a measure of uncertainty u_t and 12 month moving average of industrial production growth in period $t + k$, ie. $\text{IP-Corr}(k)=|\text{corr}(u_t, \Delta \ln IP_{t+k})|$. $\bar{\mathcal{U}}_t^y(1)$ denotes aggregate uncertainty and $\mathcal{D}(\cdot)$ represents cross-sectional standard deviation. Monthly series are aggregated by averaging observations over each larger period. The sample periods are 1960:07-2011:12 (monthly), 1961:Q3-2011:Q3 (quarterly), 1960:H2-2011:H2 (half-years), 1960-2009 (annual).

First, the estimated half life of a shock to aggregate uncertainty is 53 months. This can be compared to a common proxy for uncertainty, the VXO stock market volatility index constructed by the Chicago Board of Options Exchange (CBOE) from the prices of options contracts written on the S&P 100 Index.¹⁶ The estimated half life of a shock to stock market volatility (VXO) is 4 months. Thus, macro uncertainty is much more persistent than the most common proxy for uncertainty, a finding relevant for theories where uncertainty is a driving force of economic downturns, including those with more prolonged periods of below-trend economic growth. Second, the skewness of $\bar{U}_t^y(1)$ is similar to that for VXO, but the kurtosis of $\bar{U}_t^y(1)$ is lower than VXO. This implies that there are more extreme values in VXO, consistent with the visual inspection of the two series. Third, aggregate uncertainty is strongly countercyclical and has a contemporaneous correlation with industrial production of -0.62. Moreover, a substantial part of the comovement between aggregate uncertainty and production is attributable to uncertainty leading real activity. The maximum of IP-Corr(k) conditional on $k > 0$ is -0.67 and occurs at $k = 3$. But there is also a substantial component of the comovement in which uncertainty lags real activity. At negative values of k , the maximum of IP-Corr(k) is -0.59 and occurs at $k = -1$. By contrast, at a one year horizon (corresponding to $k = 12, 4, 2, 1$ in monthly, quarterly, semi-annual and annual data), $\bar{U}_t^y(1)$ has a much stronger correlation with future real activity than past (-0.44 versus -0.14 in monthly data), where as the opposite is true for the cross-sectional variance of firm profits and the dispersion in subjective GDP forecasts. Of course, these unconditional correlations are uninformative about the causal relation between uncertainty and real activity. All that can be said is that there is a strong coherence between uncertainty and real activity.

Uncertainty in a series is defined above as the volatility of a purely unforecastable error

¹⁶This index is available from 1986. Following Bloom (2009), we create a longer series by splicing the CBOE VXO with estimates of realized stock market volatility for the months before 1986. Specifically, from 1961:1-1986:12, the series is the standard deviation of stock returns and from 1986:1-2011:12, the series is VXO from the CBOE. We refer to this spliced version as the VXO index. The VXO series is used instead of the VIX because the VIX data does not exist before 1990. We have also constructed a new “VIX” series that splices the post 1990 VIX from CBOE with the standard deviation of stock returns for the earlier sample 1960: 1-1989:12 and found very similar results. The correlation between VXO and VIX is 0.99 over the overlapping sample.

of that series. It is potentially influenced by macro uncertainty shocks and idiosyncratic uncertainty shocks. To assess the relative importance of macro uncertainty $\overline{\mathcal{U}}_t^y(h)$ in total uncertainty (summed over all series), we compute, for each of the 132 series in the macro dataset

$$R_{j\tau}^2(h) = \frac{\text{var}_\tau(\hat{\varphi}_{j\tau}(h)\overline{\mathcal{U}}_t^y(h))}{\text{var}_\tau(\widehat{\mathcal{U}}_{jt}^y(h))}. \quad (2.14)$$

where $\hat{\varphi}_{j\tau}(h)$ is the coefficient from a regression of $\widehat{\mathcal{U}}_{jt}^y(h)$ on $\overline{\mathcal{U}}_t^y(h)$. Thus $R_{j\tau}^2(h)$ is the fraction of variation in $\widehat{\mathcal{U}}_{jt}^y(h)$ explained by macro uncertainty $\overline{\mathcal{U}}_t^y(h)$ in the subsample. The statistic is computed for $h = 1, 3$, and 12 , for the full sample, for recession months, and for non-recession months.¹⁷ The larger is $R_t^2(h) \equiv \frac{1}{N_y} \sum_{j=1}^{N_y} R_{jt}^2(h)$, the more important is macro uncertainty in explaining total uncertainty.

Table 2.2 shows that the importance of macro uncertainty grows as the forecast horizon h increases. On average across all series, the fraction of series uncertainty that is driven by common macro uncertainty is much higher for $h = 3$ and $h = 12$ than it is for $h = 1$. Table 2.2 also shows that macro uncertainty $\overline{\mathcal{U}}_t^y(h)$ accounts for a quantitatively large fraction of the variation in total uncertainty in the individual series. When the uncertainty horizon is $h = 3$ months, estimated macro uncertainty explains an average (across all series) 24%. But because there is much more variability in uncertainty in recessions, the amount explained in recessions is much larger (26%) than in non-recessions (16%). The results are similar for the $h = 12$ case. Results in the right panel of the table based on the common uncertainty factor $\overline{\mathcal{U}}_t(h)$ constructed by the method of principal components reinforce the point that macro uncertainty accounts for a larger fraction of the variation in total uncertainty during recessions.

These results show that, on average across series, macro uncertainty is quantitatively important. But there is a large amount idiosyncratic variation in uncertainty across series, as evident from the many $R_t^2(h)$ statistics that are substantially lower than unity. This is also

¹⁷Recession months are defined by National Bureau of Economic Research dates. Macro uncertainty is estimated over the full sample even when the R^2 statistics are computed over subsamples.

evident from examining additional results (not in the table) on the average R^2 in the lowest and highest quartile of series. Whereas the average of $R_{jt}^2(h)$ for $h = 3$ is 0.24, the upper and lower quartiles are 0.11 and 0.37, respectively. For an uncertainty horizon of $h = 12$ months, the three series that are most explained by macro uncertainty contemporaneously are: manufacturing and trade inventories relative to sales, housing starts (South), and housing starts (nonfarm), with R^2 equal to 0.8, 0.8, and 0.78, respectively. The three series that are least explained by macro uncertainty contemporaneously are: NAPM vendor deliveries index, CPI-U (medical care), and a measure of the number of long-run unemployed (persons unemployed 27 weeks or more). All of these have R^2 that are effectively zero.

Average R^2 From Regressions of Individual Uncertainty on Macro Uncertainty						
Average: $\bar{\mathcal{U}}^y(h) = \frac{1}{N_y} \sum_{j=1}^{N_y} \hat{\mathcal{U}}_{jt}(h)$				PC: $\bar{\mathcal{U}}^y(h) = \sum_{j=1}^{N_y} w_j \hat{\mathcal{U}}_{jt}(h)$		
h	full sample	recession	non-recession	full sample	recession	non-recession
1	0.18	0.19	0.12	0.17	0.17	0.12
2	0.22	0.24	0.15	0.22	0.24	0.15
3	0.24	0.26	0.16	0.23	0.25	0.16
4	0.25	0.26	0.17	0.23	0.24	0.16
5	0.26	0.27	0.18	0.24	0.25	0.16
6	0.27	0.28	0.19	0.25	0.26	0.16
7	0.28	0.29	0.19	0.25	0.27	0.17
8	0.29	0.30	0.20	0.25	0.27	0.17
9	0.29	0.30	0.20	0.25	0.28	0.17
10	0.29	0.31	0.21	0.25	0.28	0.17
11	0.30	0.31	0.21	0.25	0.29	0.17
12	0.30	0.31	0.21	0.25	0.29	0.17

Table 2.2: Cross-sectional averages of R^2 . Values are from regressions of $\hat{\mathcal{U}}_{jt}^y(h)$ on the benchmark (average across series) macro uncertainty measure $\bar{\mathcal{U}}_t^y(h)$ or the principal components (PC) macro uncertainty measure $\bar{\mathcal{U}}_t^y(h)$ over different subsamples. Uncertainty is estimated from the monthly, macro dataset. Recession months are defined according to the NBER Business Cycle Dating Committee. The data are monthly and span the period 1960:07-2011:12.

2.5.1 The Role of the Predictors

We have emphasized the importance of removing the predictable variation in a series so as not to attribute its fluctuations to a movement in uncertainty. How important are these predictable variations in our estimates? Our forecasting regression is

$$y_{jt+1} = \phi_j^y(L)y_{jt} + \gamma_j^F(L)\hat{\mathbf{F}}_t + \gamma_j^W(L)\mathbf{W}_t + \sigma_{jt+1}^y\varepsilon_{jt+1}.$$

The future values of our predictors \mathbf{F} and \mathbf{W} are unknown and each predictor is forecasted by an AR(4) model. As explained above, time-varying volatility in their forecast errors also contributes to h -step-ahead uncertainty in the variable y_{jt} whenever $h > 1$. Figure 2.2 plots estimated factor uncertainty $\hat{\mathcal{U}}_{kt}^F(h)$ for several estimated factors \hat{F}_{kt} that display significant stochastic volatility and that are frequently chosen as predictor variables according to the hard thresholding rule. These are, \hat{F}_{1t} (highly correlated with the stock market), \hat{F}_{2t} (highly correlated with measures of real activity such as industrial production and employment), \hat{F}_{4t} (highly correlated with measures of inflation), \hat{F}_{5t} (highly correlated with the Fama-French risk factors and bond default spreads). This figure also displays estimates of uncertainty for two predictors in W : the squared value of the first factor \hat{F}_{1t}^2 and for the first factor formed from observations X_{it}^2 , which we denote \hat{G}_{1t} . These results suggest that uncertainty in the predictor variables is an important contributor to uncertainty in the series y_{jt+h} to be forecast.

In addition to the stochastic volatility effect, the predictors directly affect the level of the forecast. An important aspect of our uncertainty measure is a forecasting model that exploits as much available information as possible to control for the economic state, so as not to erroneously attributing forecastable variations (as reflected in $\hat{\mathbf{F}}_t$ and \mathbf{W}_t) to uncertainty in series y_{jt+h} . Most popular measures of uncertainty do not take these systematic forecasting relationships into account. To examine the role that this information plays in our estimates, we re-estimate the uncertainty for each series based on the following (potentially misspecified)

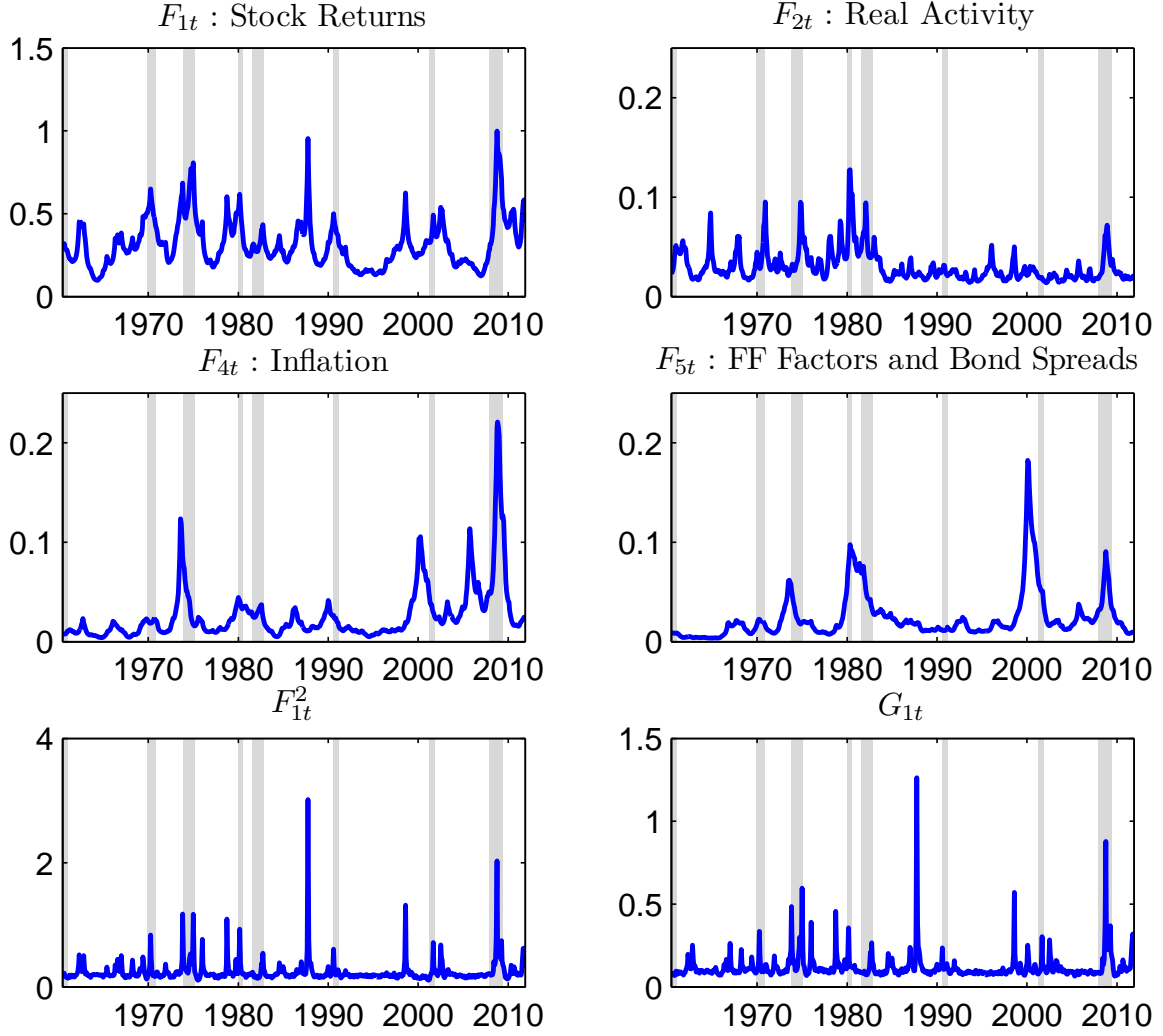


Figure 2.2: Predictor Uncertainty: This plot displays uncertainty estimates for 6 of the 14 predictors contained in the vector $Z_t \equiv (F_t', W_t')'$. F_t denotes the 12 factors estimated from X_{it} , and $W_t \equiv (F_{1t}^2, G_{1t})'$, where G_{1t} is the first factor estimated from X_{it}^2 . Titles represent the types of series which load most heavily on the factor plotted; “FF Factors” means the Fama-French factors (HML, SMB, UMD). The data are monthly and span the period 1960:07-2011:12.

simple model with constant conditional mean:

$$y_{jt+1} = \mu + \tilde{\sigma}_{jt+1} \tilde{\varepsilon}_{jt+1}. \quad (2.15)$$

Figure 2.3 plots the resulting estimates of one-step ahead uncertainty $\overline{\mathcal{U}}_{jt}^y(1)$ using this possibly misspecified model and compares it to the corresponding estimates using the full set of chosen predictors (chosen using the hard thresholding rule described above), for several key series in our dataset: total industrial production, employment in manufacturing, non-farm housing starts, consumer expectations, M2, CPI-inflation, the ten-year/federal funds term spread, and the commercial paper/federal funds rate spread. Figure 2.3 shows that there is substantial heterogeneity in the time-varying uncertainty estimates across series, suggesting that a good deal of uncertainty is series-specific. But Figure 2.3 also shows that the estimates of uncertainty in these series are significantly influenced by whether or not the forecastable variation is removed before computing uncertainty: when it is removed, the estimates of uncertainty tend to be lower, much so in some cases. Specifically, uncertainty in each of the eight variables shown in this figure is estimated to be lower during the 2007-09 recession when predictive content is removed than when not, especially for industrial production, employment, and the two interest rate spreads. The difference over time between the two estimates for these variables is quite pronounced in some periods, suggesting that much of the variation in these series is predictable and should not be attributed to uncertainty.¹⁸

Since stock market volatility is the most commonly used proxy for uncertainty, we further examine in Figure 2.4 how estimates of stock market uncertainty are affected by whether or not the purely forecastable variation in the stock market is removed before computing uncertainty. This figure compares (i) the estimate of uncertainty in the log difference of the *S&P* 500 index for a case where the conditional mean is assumed constant, implying as in

¹⁸We have also re-estimated common macro uncertainty, $\overline{\mathcal{U}}_t^y(h)$ without removing predictable fluctuations. The spikes appear larger than the base case that removes the forecastable component in each series before computing uncertainty. This is especially true for the $h = 1$ case, where presumably the predictive information is most valuable.

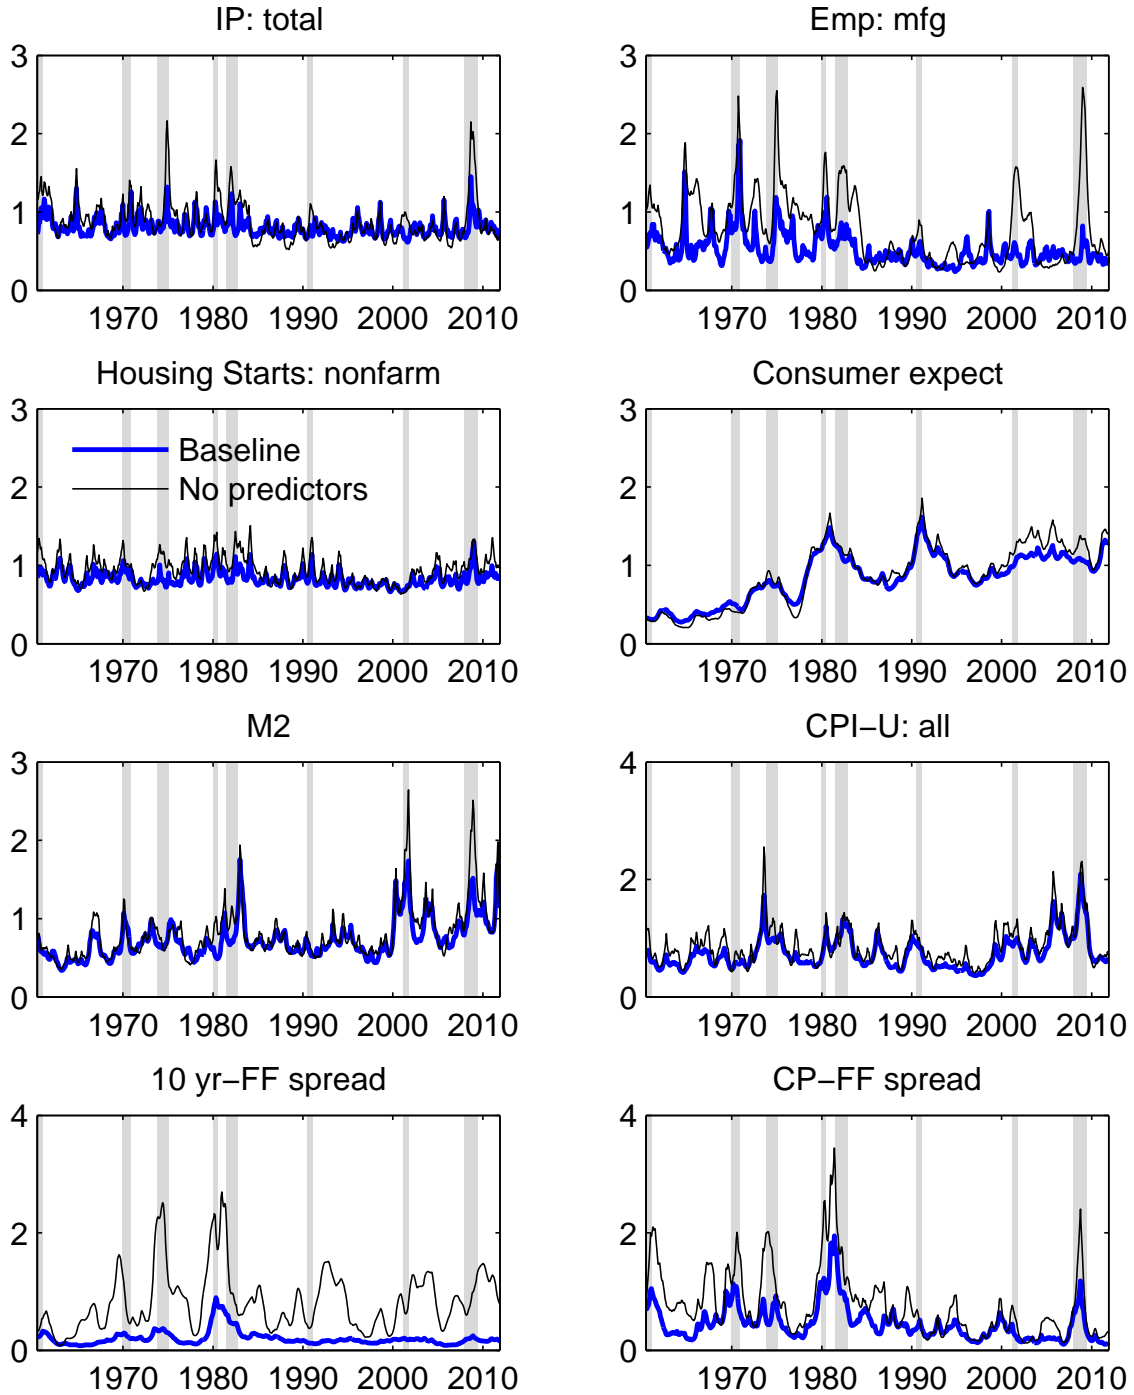


Figure 2.3: The Role of Predictors: These plots display two estimates of $\mathcal{U}_{jt}^y(1)$ for several key series in our data set. The first is constructed using the full set of predictor variables (“Baseline”); the second is constructed using no predictors (“No predictors”). The data are monthly and span the period 1960:07-2011:12.

(2.15) that no predictable variation is removed, with (ii) a case in which only autoregressive terms are included to forecast the stock market, as in

$$y_{jt+1} = \tilde{\phi}_j(L)y_{jt} + \tilde{\sigma}_{jt+1}\tilde{\varepsilon}_{jt+1},$$

with (iii) a case in which all selected factors (using the hard thresholding rule) estimated from the combined macro and financial dataset with 279 indicators are used as predictors. Notice that the first case (constant conditional mean) is most akin to estimates of stock market volatility such as the VXO index and discussed further below. We emphasize that stock market volatility measures do not purge movements in the stock market of its predictable component and are therefore estimates of conditional volatility, not uncertainty. Of course, if there were no predictable component in the stock market, these two estimates would coincide. But Figure 2.4 shows that there is a substantial predictable component in the log change in the *S&P* price index, which, once removed, makes a quantitatively large difference in the estimated amount of uncertainty over time.¹⁹ Uncertainty in the stock market is substantially lower in every episode when these forecastable fluctuations are removed compared to when they are not, and is dramatically lower in the recession of 2007-09 compared to what is indicated by ex-post conditional stock market volatility.

If we examine more closely our measure of stock market uncertainty, (given by the baseline estimate in Figure 2.4) and compare it to macro uncertainty $\overline{\mathcal{U}}_t^y$ (Figure 2.1), we see there are important differences over time in the two series. In particular, there are many (more) large spikes in stock market uncertainty that are not present for macro uncertainty. Unlike macro uncertainty, several of the spikes in financial uncertainty occur outside of recessions. Because stock market volatility is arguably the most common proxy for uncertainty, we further examine the distinction between uncertainty and stock market volatility in the next

¹⁹Evidence for predictability of stock returns is not hard to find. [Cochrane \(1994\)](#) found an important transitory component in stock prices. [Ludvigson and Ng \(2007\)](#) found substantial predictive information for excess stock market returns in the factors formed from the financial dataset \mathbf{X}^f . For more general surveys of the predictable variation in stock market returns, see [Cochrane \(2005\)](#) and [Lettau and Ludvigson \(2004\)](#).

section.




Figure 2.4: Uncertainty in the S&P 500 Index. These plots show estimates of $\mathcal{U}_{SP500,t}^y(1)$ for the S&P 500 Index based on three different forecasting models. “No Predictors” indicates that no predictors were used, “AR only” indicates that only a fourth-order autoregressive model was used to generate forecast errors, and “Baseline” indicates that the full set of predictor variables was used to generate forecast errors. The data are monthly and span the period 1960:07-2011:12.

2.5.2 Uncertainty Versus Stock Market Volatility

In an influential paper, [Bloom \(2009\)](#) emphasizes a measure of stock market volatility as a proxy of uncertainty.²⁰ This measure is primarily based on the VXO Index. In this

²⁰A number of other papers also use stock market volatility to proxy for uncertainty; these include [Romer \(1990\)](#), [Leahy and Whited \(1996\)](#), [Hassler \(2001\)](#), [Bloom et al. \(2007\)](#), [Greasly and Madsen \(2006\)](#), [Gilchrist](#)

subsection we compare our macro uncertainty estimates with stock market volatility as a proxy for uncertainty. We update this stock market volatility series to include more recent observations, and plot it along with our estimated macro uncertainty $\overline{U}_t^y(h)$ for $h = 1$ in Figure 2.5. To construct his benchmark measure of uncertainty “shocks” (plausibly exogenous variation in his proxy of uncertainty), Bloom selects 17 dates (listed in his Table A.1) which are associated with stock market volatility in excess of 1.65 standard deviations above its HP-detrended mean. These 17 dates are marked by vertical lines in the figure. As emphasized above and seen again in Figure 2.5, $\overline{U}_t^y(1)$ exceeds 1.65 standard deviations above its unconditional mean in only three episodes, suggesting far fewer episodes of uncertainty than that indicated by these 17 uncertainty dates.²¹

While $\overline{U}_t^y(1)$ is positively correlated with the VXO Index, with a correlation coefficient around 0.5, the VXO Index is itself substantially more volatile than $\overline{U}_t^y(1)$, with many sharp peaks that are not correspondingly reflected by the macro uncertainty measure. For example, the large spike in October 1987 reflects “Black Monday,” which occurred on the 19th of the month when stock markets experienced their largest single-day percentage decline in recorded history. While this may accurately reflect the sudden increase in financial market volatility that occurred on that date, our measure of macroeconomic uncertainty barely increases at all. Indeed, it is difficult to imagine that the level of macro uncertainty in the economy in October 1987 (not even a recession year) was on par with the recent financial crisis. Nevertheless, when the VXO index is interpreted as a proxy for uncertainty, this is precisely what is implied. Other important episodes where the two measures disagree include the recessionary period from 1980-1982, where our measure of uncertainty was high but the VXO index was comparatively low, and the stock market boom and bust of the late 1990s and early 2000s, where the VXO index was high but uncertainty was low.

et al. (2014), and Basu and Bundick (2012).

²¹Bloom (2009) counts uncertainty episodes by the number of times the stock market volatility index exceeds 1.65 standard deviations above its Hodrick-Prescott filtered trend, rather than its unconditional mean. If we do the same for $\overline{U}_t^y(1)$, we find 5 episodes of heightened uncertainty: one in the early mid 1970s (1973:09 and 1974:11), one during the twin recessions in the early 1980s (1980:02 and 1982:02), 1990:01, 2001:10, and 2008:07.




Figure 2.5: Stock Market Implied Volatility and Uncertainty: This plot shows $\bar{U}_t^y(1)$ and the VXO index, expressed in standardized units. The vertical lines correspond to the 17 dates in Bloom (2009) Table A.1, which correspond to dates when the VXO index exceeds 1.65 standard deviations above its HP (Hodrick and Prescott, 1997) filtered mean. The horizontal line corresponds to 1.65 standard deviations above the unconditional mean of each series (which has been normalized to zero). The data are monthly and span 1960:07-2011:12.

2.5.3 Macro Uncertainty and Macroeconomic Dynamics

Existing empirical research on uncertainty has often found important dynamic relationships between real activity and various uncertainty proxies. In particular, these proxies are countercyclical and VAR estimates suggest that they have a large impact on output and employment in the months after an innovation in these measures. A key result is that a rise in some proxies (notably stock market volatility) at first depresses real activity and then increases it, leading to an over-shoot of its long-run level, consistent with the predictions of some theoretical models on uncertainty as a driving force of macroeconomic fluctuations.

We now use VARs to investigate the dynamic responses of key macro variables to innovations in our uncertainty measures and compare them to the responses to innovations in the VXO index as a proxy for uncertainty. For brevity in discussing the results, we will often refer to these innovations to uncertainty or stock market volatility (in the case of the VXO index) as “shocks.” As is the case of all VAR analyses, the impulse responses and variance decompositions depend on the identification scheme, which in our case is based on the ordering of the variables.

A question arises as to which variables to include in the VAR. As a starting point, we choose a macro VAR similar to that studied in Christiano, Eichenbaum and Evans (2005, CEE hereafter). This VAR affords the advantage of containing a set of variables whose dynamic relationships have been the focus of extensive macroeconomic research. Since CEE use quarterly data and we use monthly data, we do not use exactly the same VAR, but instead include similar variables so as to roughly cover the same sources of variation in the economy.²² We estimate impulse responses from a eleven-variable VAR, hereafter referred

²²Specifically, monthly industrial production and the PCE deflator are substituted for quarterly Gross Domestic Product GDP and its deflator, hours is used instead of labor productivity, average hourly earnings is for the manufacturing sector only because the aggregate measure does not go back to 1960, and the S&P 500 stock market index is substituted for quarterly corporate profits.

to as VAR-11. The ordering mimics that of CEE:

$$\begin{bmatrix} \log(\text{real IP}) \\ \log(\text{employment}) \\ \log(\text{real consumption}) \\ \log(\text{PCE deflator}) \\ \log(\text{real new orders}) \\ \log(\text{real wage}) \\ \text{hours} \\ \text{federal funds rate} \\ \log(\text{S\&P 500 Index}) \\ \text{growth rate of M2} \\ \textit{uncertainty} \end{bmatrix} \quad (\text{VAR-11})$$

Four versions of VARs-11 with twelve lags are considered with *uncertainty* taken to be either $\bar{U}_t^y(1)$, $\bar{U}_t^y(3)$, $\bar{U}_t^y(12)$, or the VXO Index. The main difference from the CEE VAR is the inclusion of a stock price index and uncertainty. It is important to include the stock market index for understanding the dynamics of uncertainty since it is natural to expect the two variables to be dynamically related. In all cases, we place the measure of uncertainty last in the VAR.

The shocks to which dynamic responses are traced are identified using a Cholesky decomposition, with the same timing assumptions made in CEE that allows identification of federal funds rate shocks.²³

In addition to [VAR-11](#), it is also of interest to compare the dynamic correlations of our uncertainty measures with common uncertainty proxies using a VAR that has been previously employed in the uncertainty literature. To do so, we estimate impulse responses

²³We have confirmed that the dynamic responses of the non-uncertainty variables to a federal funds rate shock (interpreted by CEE as a monetary policy shock) in a VAR that does not include any uncertainty measure are qualitatively and quantitatively very similar to those reported in CEE. These results are available upon request.

from a eight-variable model as in [Bloom \(2009\)](#), hereafter referred to as [VAR-8](#):

$$\begin{bmatrix} \log(\text{S\&P 500 Index}) \\ \textit{uncertainty} \\ \text{federal funds rate} \\ \log(\text{wages}) \\ \log(\text{CPI}) \\ \text{hours} \\ \log(\text{employment}) \\ \log(\text{industrial production}) \end{bmatrix}. \quad (\text{VAR-8})$$

Following [Bloom \(2009\)](#), [VAR-8](#) uses twelve lags of industrial production, wages, hours. Unlike [VAR-11](#), [VAR-8](#) uses employment for the manufacturing sector *only*. [Bloom \(2009\)](#) considers a 15-point shock to the error in the VXO equation. This amounts to approximately 4 standard deviations of the identified error. We record responses to 4 standard deviation shocks in $\bar{\mathcal{U}}_t^y(h)$, so the magnitudes are comparable with those of VXO shocks. However, we make one departure from the estimates in [Bloom \(2009\)](#). We do not detrend any variables using the filter of [Hodrick and Prescott \(1997\)](#), while Bloom did so for every series except the VXO index. Because the HP filter uses information over the entire sample, it is difficult to interpret the timing of an observation.²⁴

Figure 2.6 shows the dynamic responses of output and employment in [VAR-11](#). Shocks to $\bar{\mathcal{U}}_t^y(h)$ sharply reduce production and employment, with the effects persisting well past the 60 month horizon depicted. The last row of this figure compares the responses when the VXO index is used as a proxy for uncertainty. Both the magnitude and the persistence of the responses of production and employment are much smaller. The responses to $\bar{\mathcal{U}}_t^y(h)$ are far more protracted than those to the VXO Index, which underscores the greater persistence

²⁴Results using HP filtered data and the original Bloom VAR are reported in the on-line supplementary material file for this paper.

of these measures as compared to popular uncertainty proxies. Indeed, the response of employment to a VXO disturbance is barely statistically different from zero shortly after the shock and outright insignificant at other horizons. The response of production to a VXO shock is also only marginally different from zero for the first 3 months, becoming zero thereafter. An important difference in these results from those reported in [Bloom \(2009\)](#) is that shocks to any of these measures (including VXO) do not generate a statistically significant “volatility overshoot,” namely, the rebound in real activity following the initial decline after a positive uncertainty shock. This finding echoes those in [Bachmann and Bayer \(2011\)](#). Unlike the findings in [Bachmann and Bayer \(2011\)](#), however, the short-run (within 10 months) responses to our uncertainty shocks are sizable.

Figure [2.7](#) shows the dynamic responses of output and employment in [VAR-8](#). The responses of these variables, both in terms of magnitude and persistence, to the macro uncertainty measures $\bar{\mathcal{U}}_t^y(h)$ are similar to those reported in Figure [2.6](#) using [VAR-11](#). Disturbances to the VXO index appear to have larger and somewhat more persistent effects in [VAR-8](#) than in [VAR-11](#). But the responses to VXO shocks even in this VAR are not as large or persistent as those to innovations in macro uncertainty $\bar{\mathcal{U}}_t^y(h)$. Again, there is no clear evidence of a volatility overshoot in response to any of the uncertainty measures, including VXO. The overshoot found by [Bloom \(2009\)](#) appears to be sensitive to whether the VXO data are HP filtered.²⁵

To study the quantitative importance of uncertainty shocks for macroeconomic fluctuations, Table [2.3](#) reports forecast error variance decomposition for production, employment and hours and compares them with the decompositions when VXO is used instead as the proxy for uncertainty in the [VAR-11](#). We use k here to distinguish the VAR forecast horizon from the uncertainty forecast horizon h . The table shows the fraction of the VAR forecast error variance that is attributable to common macro uncertainty shocks in $\bar{\mathcal{U}}_t^y(h)$ over several horizons, including the horizon k for which shocks to the uncertainty measure $\bar{\mathcal{U}}_t^y(1)$ or VXO

²⁵After a careful inspection of the code kindly provided by Bloom, we find that contrary to a statement in the paper, [Bloom \(2009\)](#) HP filters all data in the VAR for these impulse responses *except* the VXO index.

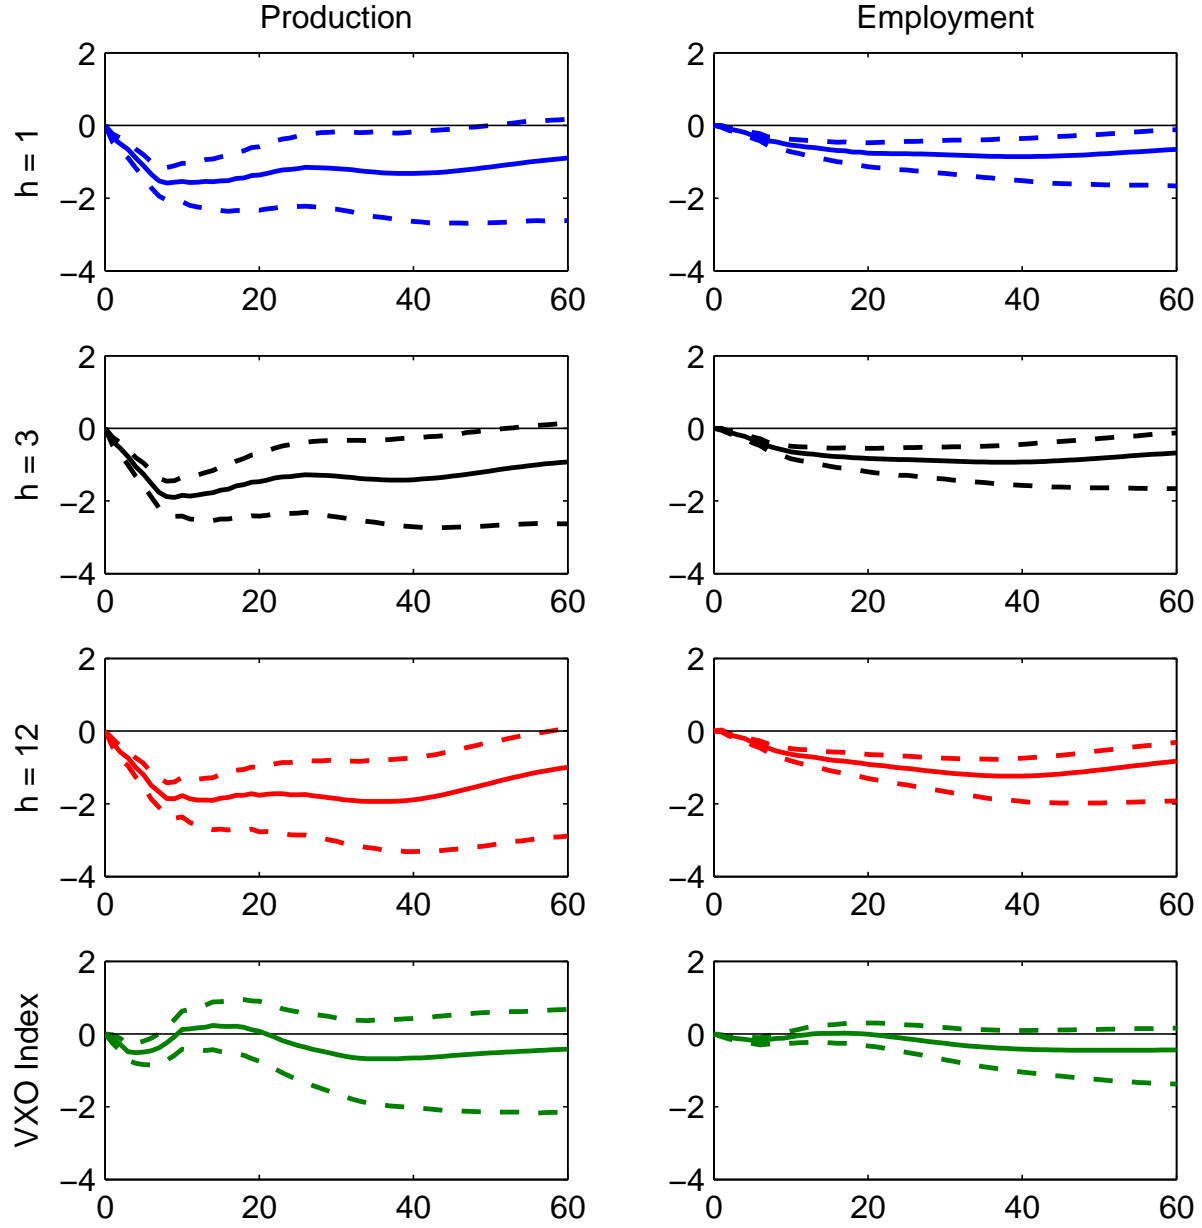


Figure 2.6: Impulse response of production and employment from estimation of [VAR-11](#) using $\bar{U}_t^y(h)$ or VXO as uncertainty. Dashed lines show 68% standard error bands. The data are monthly and span the period 1960:07 to 2011:12.

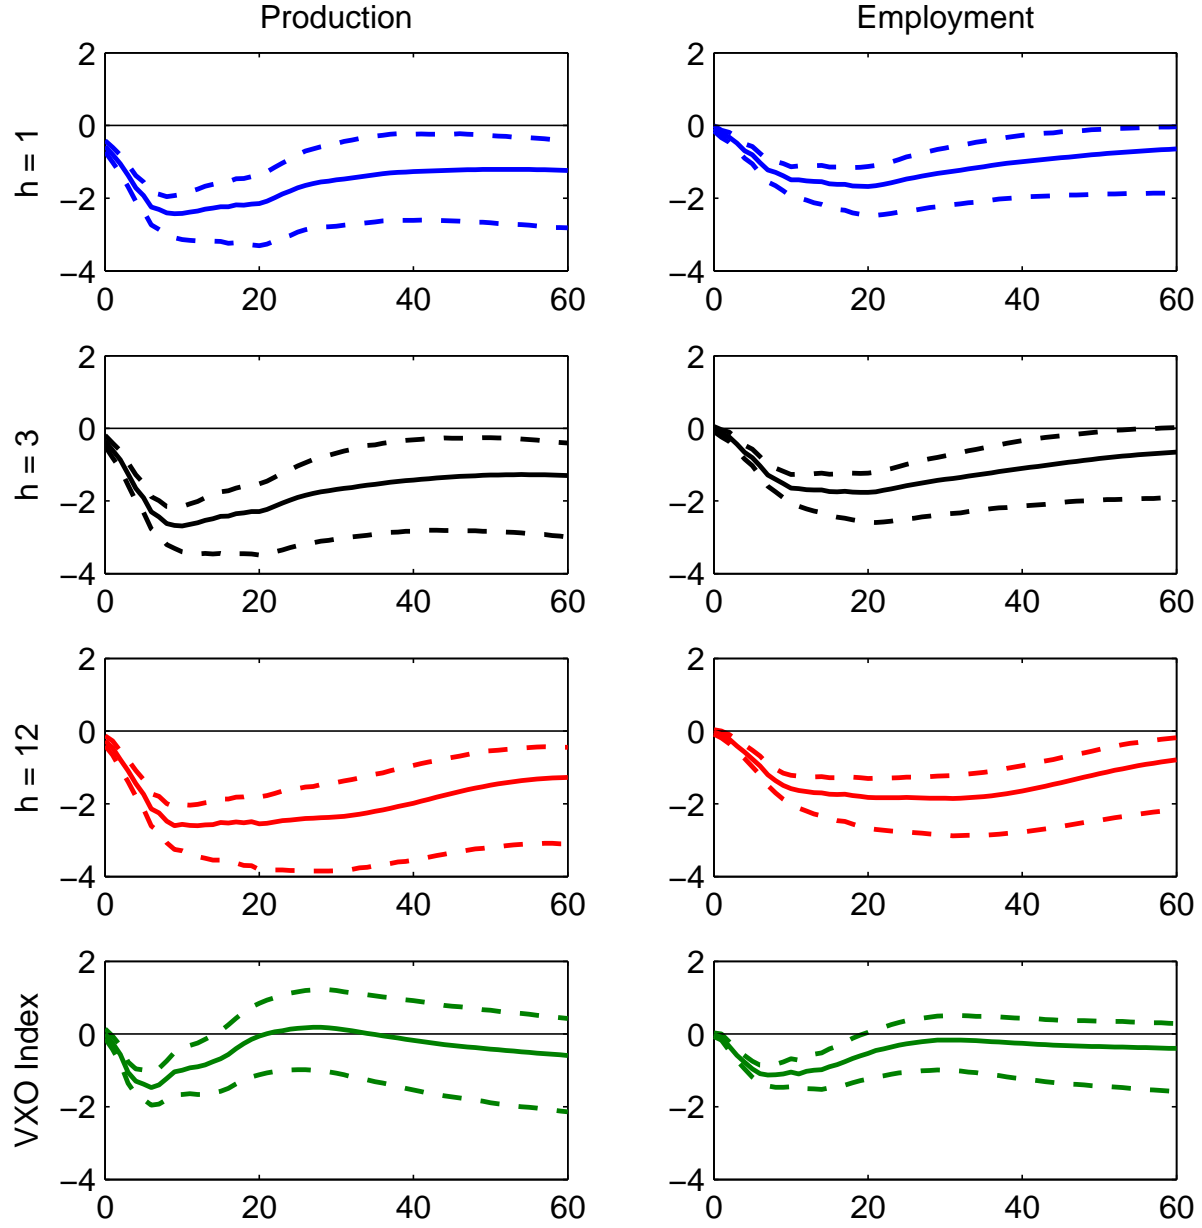


Figure 2.7: Impulse response of production and employment from estimation of VAR-8 using $\bar{U}_t^y(h)$ or VXO as uncertainty. Dashed lines show 68% standard error bands. The data are monthly and span the period 1960:07 to 2011:12.

are associated with the greatest fraction of VAR forecast error variance (denoted $k = \max$ in the table). The table also reports the fraction of variation attributable to the federal funds rate, which we discuss below.

From Table 2.3 we can see that uncertainty shocks are associated with much larger fractions of real activity than are VXO shocks. Shocks to $\bar{U}_t^y(12)$, for example, are associated with a maximum of 29% of the forecast error variance in production, 31% of the forecast error variance in employment, and 12% of the forecast error variance in hours. By contrast, the corresponding numbers for VXO shocks are 6.9%, 7.6%, and 2.3%, respectively. Thus, uncertainty shocks are associated with over four times the variation in production and employment and over five times the variation in hours compared to VXO shocks.

To put these results in perspective, Table 2.3 also reports the fraction of variation in these variables that is attributable to monetary policy shocks, identified here following CEE by a shock in VAR-11 to the federal funds rate. In the VAR-11 where $\bar{U}_t^y(12)$ is included as the measure of uncertainty, shocks to the federal funds rate are associated with a maximum of 29% of the forecast error variance in production, 32% of the forecast error variance in employment, and 10% of the forecast error variance in hours. These numbers are almost identical to the fraction explained by shocks to $\bar{U}_t^y(12)$. This finding suggests that the dynamic correlation of uncertainty with the real economy may be quantitatively as important as it is for monetary policy shocks.

We can use the same variance decompositions to ask how much of uncertainty variation is associated with variation in innovations of the other variables in the system. These results are not reported in the table, but we discuss a few of them here. At the $k = \infty$ horizon, we find that stock return innovations are associated with the largest fraction of variation in $\bar{U}_t^y(12)$, equal to 15.26%, followed by price level innovations (11.9%) and innovations to industrial production (9.56%). These numbers are roughly of the same order of magnitude as those for the fraction of forecast error variance in production growth explained by $\bar{U}_t^y(12)$ for $k = \infty$ (equal to 15.75%). These variance decompositions are of course specific to the

Relative Importance of Uncertainty v.s. FFR in VAR-11

Fraction Variation in Production (%)								
Explained by:	$\bar{\mathcal{U}}_t^y(1)$	FFR	$\bar{\mathcal{U}}_t^y(3)$	FFR	$\bar{\mathcal{U}}_t^y(12)$	FFR	VXO	FFR
$k = 3$	1.78	0.06	2.08	0.04	2.13	0.02	0.48	0.01
$k = 12$	11.29	5.86	15.79	5.27	15.22	4.00	0.91	7.17
$k = \infty$	7.87	33.67	8.79	31.39	15.76	28.96	6.93	39.07
max k	174	∞	171	∞	174	∞	184	∞
$k = \max$	17.02	33.67	20.86	31.39	28.54	28.96	6.93	39.07
Fraction Variation in Employment (%):								
Explained by:	$\bar{\mathcal{U}}_t^y(1)$	FFR	$\bar{\mathcal{U}}_t^y(3)$	FFR	$\bar{\mathcal{U}}_t^y(12)$	FFR	VXO	FFR
$k = 3$	0.90	0.06	0.98	0.03	0.86	0.01	1.06	0.02
$k = 12$	9.15	6.99	13.23	6.33	13.08	4.87	1.11	8.26
$k = \infty$	6.66	36.02	7.51	33.14	14.25	31.89	7.64	39.47
max k	105	185	106	190	107	357	184	148
$k = \max$	16.40	41.30	20.06	39.35	31.00	34.83	7.64	52.74
Fraction of Variation in Hours (%):								
Explained by:	$\bar{\mathcal{U}}_t^y(1)$	FFR	$\bar{\mathcal{U}}_t^y(3)$	FFR	$\bar{\mathcal{U}}_t^y(12)$	FFR	VXO	FFR
$k = 3$	1.76	0.44	1.88	0.48	1.26	0.56	0.12	0.72
$k = 12$	8.11	4.58	11.36	4.30	10.53	3.54	1.16	6.36
$k = \infty$	7.38	12.92	8.98	12.08	11.93	9.79	2.15	17.21
max k	21	∞	16	∞	37	∞	43	∞
$k = \max$	9.21	12.92	11.96	12.08	12.34	9.79	2.32	17.21

Table 2.3: Decomposition of variance in production, employment and hours due to either uncertainty or the federal funds rate in VAR-11. The VAR uses variables in the following order: log(industrial production), log(employment), log(real consumption), log(implicit consumption deflator), log(real value new orders, consumption and non-defense capital goods), log(real wage), hours, federal funds rate (FFR), log(S&P 500 Index), growth rate of M2, and *uncertainty*, where the latter is either $\bar{\mathcal{U}}_t^y(h)$ or the VXO Index. We estimate separate VARs in which *uncertainty* is either one of $\bar{\mathcal{U}}_t^y(h)$, $h = 1, 3, 12$ or the VXO index. Each panel shows the fraction of forecast-error variance of the variable given in the panel title at VAR forecast horizon k that is explained by the uncertainty measure, as named in the column, or the FFR for that VAR. The row denoted “max k ” gives the horizon k for which the uncertainty variable named in the column explains the maximum fraction of forecast error variance. The row denoted “ $k =$ ”max gives the fraction of forecast error variance explained at max k . Real variables are obtained by dividing nominal values by the PCE deflator. The data are monthly and span the period 1960:07-2011:12.

ordering of the variables used in the analysis. But as uncertainty is placed *last* in the VAR, the effects of uncertainty shocks on the other variables in the system are measured after we have removed all the variation in uncertainty that is attributable to shocks to the other endogenous variables in the system. That the effects of uncertainty shocks are still non-trivial is consistent with the view that uncertainty has important implications for economic activity. These variance decomposition results are similar if we instead use VARs that include both VXO and our uncertainty measures $\bar{U}_t^y(h)$. From such VARs, we find that the big driver of VXO are shocks to VXO, not uncertainty. This reinforces the conclusion that stock market volatility is driven largely by shocks other than those to broad-based economic uncertainty, suggesting researchers should be cautious when using this measure as a proxy for uncertainty.

We have reported results only for the base-case estimates described above. Appendix B provides additional results designed to check the sensitivity of our results to various assumptions made above. These exercises are based on (i) alternative weights used to aggregate individual uncertainty series; (ii) alternative location statistics of stochastic volatility to construct individual uncertainty series; (iii) alternative conditioning information based on recursive (out-of-sample) forecasts to construct diffusion index forecasts (iv) alternative measures of volatility of individual series such as GARCH and EGARCH.²⁶ The key findings are qualitatively and quantitatively similar to the ones reported here. We note one finding in particular, namely that the results above are not sensitive to whether we use out-of-sample (recursive) or in-sample forecasts; indeed the correlation between the resulting uncertainty measures is 0.98.²⁷

²⁶Results based on the GARCH/EGARCH estimates indicate the number and timing of big uncertainty episodes, as well as the persistence of uncertainty, is very similar to what is found using our base-case measure of macro uncertainty. What is different is the real effect of uncertainty innovations from a VAR, once orthogonalized shocks are analyzed. This is to be expected because GARCH type models (unlike stochastic volatility) have a shock to the second moment that is not independent of the first moment, a structure inconsistent with the assumptions of an independent uncertainty shock presumed in the uncertainty literature. Using a GARCH-based uncertainty index thus creates additional identification problems that are beyond the scope of this paper.

²⁷Note also that, in the recursive forecast estimation the parameters of the forecasting relation change every period, so this speaks directly to the question of the role played by parameter stability in our estimates, suggesting that parameter instability is not important in our FAVAR.

2.5.4 Comparison with Measures of Dispersion

This subsection compares the time-series behavior of $\overline{U}_t^y(h)$ with four cross-sectional uncertainty proxies studied by Bloom (2009). These are:

1. The cross-sectional dispersion of firm stock returns. This is defined as the within-month cross-sectional standard deviation of stock returns for firms with at least 500 months of data in the Center for Research in Securities Prices (CRSP) stock-returns file. The series is also linearly detrended over our sample period.
2. The cross-sectional dispersion of firm profit growth. Profit growth rates are normalized by average sales on a monthly basis, so that this measure captures the quarterly cross-sectional standard deviation profits. We formulate a year-over-year version to minimize seasonal variation equal to $\frac{\text{profits}_{it} - \text{profits}_{it-4}}{0.5(\text{sales}_{it} + \text{sales}_{it-4})}$, where $i = 1, 2, \dots, N_t$ indexes the firms and N_t denotes the total number of firms observed in month t . The sample is restricted to firms with at least 150 quarters of data in the Compustat (North America) database.
3. The cross-sectional dispersion of GDP forecasts from the Philadelphia Federal Reserve Bank's biannual Livingston Survey. This is defined as the biannual cross-sectional standard deviation of forecasts of nominal GDP one year ahead. The series is also linearly detrended over our sample period.
4. The cross-sectional dispersion of industry-level total factor productivity (TFP). This is defined as the annual cross-sectional standard deviation of TFP growth rates within SIC 4-digit manufacturing industries, calculated using the five-factor TFP growth data computed by Bartelsman, Becker, and Marvakov as a part of the NBER-CES Manufacturing Industry Database (<http://www.nber.org/data/nbprod2005.html>).²⁸

²⁸There is a jump in the 1997 industry TFP dispersion measure that occurs purely because of a move from NAICS to SIC industry classification codes. We therefore drop this year and interpolate to obtain the continuous panel.

These updated series, along with $\bar{U}_t^y(1)$ are displayed in Figure 2.8. As was true in the case of stock market volatility in the previous subsection, these measures exhibit quite different behavior from macroeconomic uncertainty. Stock return dispersion tells a story roughly similar to the VXO Index, with a particularly large increase in uncertainty leading up to the 2001 recession that is not present in our measure of macro uncertainty. Firm profit dispersion actually suggests a relatively low level of uncertainty during the 1980-82 recessions when macro uncertainty was high, with a sharp increase towards the end of the 1982 recession, by which time macro uncertainty had declined. GDP forecast dispersion points to a level of uncertainty during each of the 1969-70 and 1990 recessions which is on par with the level of uncertainty during the 2007-09 recession. Again, this contrasts with macro uncertainty which is at a record high in the 2007-09 recession but was not high in the previous episodes. Industry TFP dispersion shows almost no increase in uncertainty during the 1980-82 recessions and displays the largest increase during the recent financial crisis.

It is instructive to consider the different statistical properties of these dispersion measures as they compare to those for the estimated aggregate uncertainty index. Table 2.1 provides the statistics. To match the frequency of the dispersion measure, we aggregate our monthly series $\bar{U}_t^y(h)$ using averages over the desired period.

The statistics using these proxies for uncertainty paint a similar picture to that obtained using the VXO Index. In particular, the responses of $\bar{U}_t^y(1)$ to its own shock from an autoregression are far more prolonged than those of the dispersion proxies. For example, the response of the dispersion in firm-level stock returns to its own shock has a half-life of 1.9 months, compared to 52.5 months for $\hat{U}_{jt}^y(1)$.

We also consider impulse responses of production and employment for the eleven-variable VAR, but using these measures of dispersion as the proxy for uncertainty. These results are reported in Figure 2.9 and can be summarized as follows. The dynamic responses using dispersions to proxy for uncertainty do not in general display the intuitive pattern that production and employment should fall as a result of an uncertainty shock. Production falls

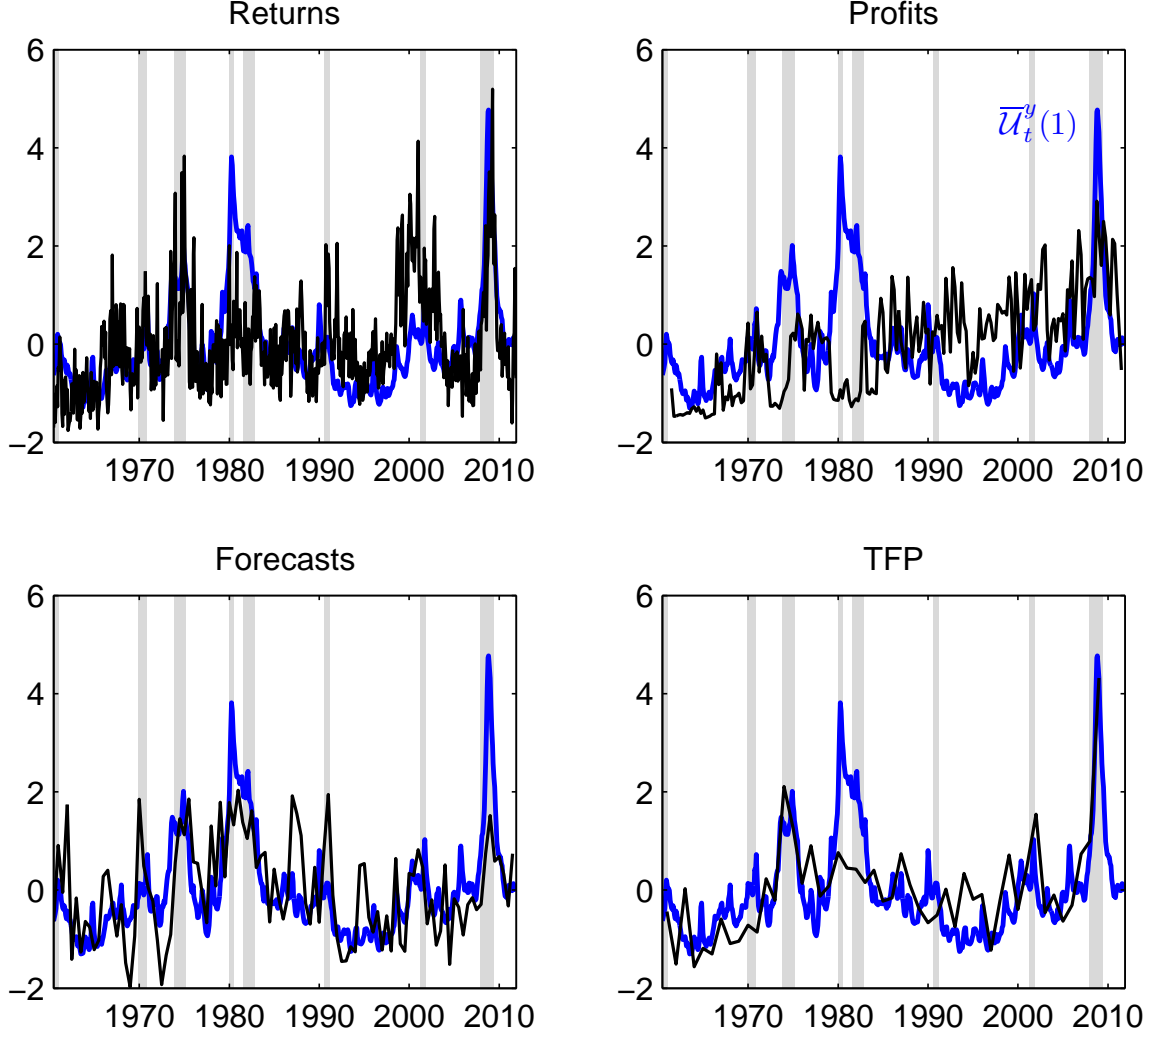


Figure 2.8: Cross-sectional Dispersion and Uncertainty: This plot shows $\bar{u}_t^y(1)$ and four dispersion-based proxies, expressed in standardized units. The proxies are (in clockwise order from the northwest panel) the cross-sectional standard deviation of: monthly firm stock returns (CRSP), quarterly firm profit growth (Compustat), yearly SIC 4-digit industry total factor productivity growth (NBER-CES Manufacturing Industry Database), and half-yearly GDP forecasts (Livingston Survey). The sample for each dataset is the largest available that overlaps with our uncertainty estimates: 1960:07-2011:12 (monthly), 1961:Q3-2011:Q3 (quarterly), 1960:H2-2011:H2 (half-years), 1960-2009 (annual).

the most on impact in response to shocks to the cross-sectional dispersion in industry-level TFP, but the response of employment is more muted. In the case of stock return dispersion, we see no statistically significant response in production or employment to an innovation. Shocks to the dispersion in firm profits lead to an *increase* in production and employment, as do shocks to the cross-sectional dispersion in subjective GDP forecasts.

Overall, these results show that, like the VXO proxy, increases in measures of cross-sectional dispersion do not necessarily coincide with increases in broad-based macro uncertainty, where the latter is associated with a large and persistent decline in real activity. Like stock market volatility over time, measures of dispersion may vary for many reasons that are unrelated to broad-based macroeconomic uncertainty.

2.6 Results: Firm-Level Common Uncertainty

In this section we turn from our analysis of common macroeconomic uncertainty to examine common variation in uncertainty at the firm level. Rather than studying uncertainty across many different variables, we now study uncertainty on the same variable across many different firms. Specifically, we measure uncertainty in the profit growth of individual firms. For the firm-level dataset, the unit of observation is the change in firm pre-tax profits $P_{i,t}$, normalized by a two-period moving average of sales, $S_{i,t}$, following Bloom (2009). Given the seasonality in this series, we instead form a year-over-year version of this measure, as detailed in the data appendix. After converting to a balanced panel, we are left with 155 firms from 1970:Q1-2011:Q2 without missing values.²⁹ With data transformations and lags in the FAVAR, we are left with uncertainty estimates for 1970:Q3-2011:Q2. For each firm, the series to be forecast is normalized pretax profits, so again $y_{it} = X_{it}$. For the firm-level results, as for the macro results, we form forecasting factors F_t from the panel $\{X_{it}\}_{i=1}^{N_{xp}}$, as well as $\{X_{it}^2\}_{i=1}^{N_{xp}}$.

²⁹A limitation with Compustat data is that its coverage is restricted to large publicly traded firms. The Census Bureau's ASM data are more comprehensive, but limited to annual observations. Similarly, (industry level) total factor productivity may be preferred over profits as the source of uncertainty, but these industry level data eliminate much of the uncertainty at the firm level (Schaal, 2011).

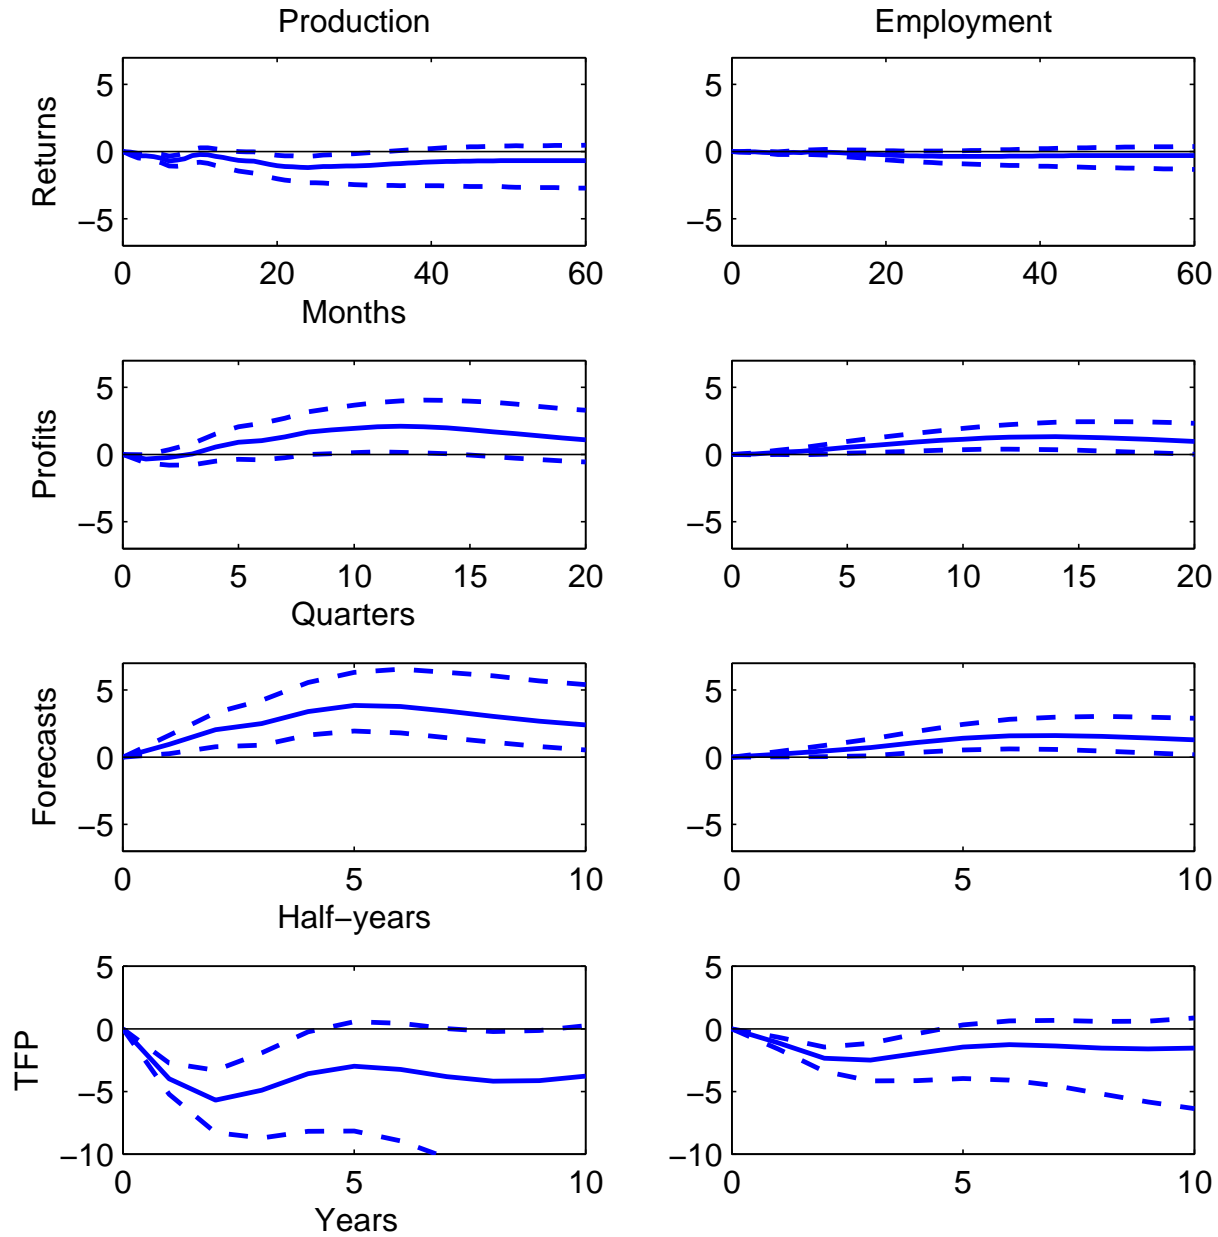


Figure 2.9: Impulse response of production and employment from estimation of VAR-11 using four dispersion measures \mathcal{D}_t as uncertainty: (i) “Returns” is the cross-sectional standard deviation of firm stock returns; (ii) “Profits” is the cross-sectional standard deviation of firm profits; (iii) “Forecasts” is the cross-sectional standard deviation of GDP forecasts from the Livingston Survey; (iv) “TFP” is the cross-sectional standard deviation of industry-level total factor productivity. Dashed lines show 68% standard error bands. The sample for each dataset is the largest available that overlaps with our uncertainty estimates: 1960:07-2011:12 (monthly), 1961:Q3-2011:Q3 (quarterly), 1960:H2-2011:H2 (half-years), 1960-2009 (annual).

where $N_{xp} = 155$, the number of cross-sectional firm-level observations. We find evidence of two factors in $\{X_{it}\}_{i=1}^{N_{xp}}$ and one factor in $\{X_{it}^2\}_{i=1}^{N_{xp}}$. The \mathbf{W}_t vector of additional predictors includes the macro factors estimated from the macro data set. As before, a conservative t test is used to include only the predictors that are statistically significant.

One important consideration that is relevant to this microeconomic context is the construction of our panel. Since we need a reasonable number of time series observations to estimate the stochastic volatility processes, we require that the panel be balanced. This leads us to drop about 400 firms per quarter on average. In particular, many of the firms operating towards the beginning of our sample are excluded, because they do not survive until 2011:Q2. This eliminates a large fraction of the cross-sectional variation before 1995. Because of this survivorship bias, it is difficult to conclude that our estimated aggregate firm-level uncertainty measure represents a comprehensive measure of the uncertainty facing firms since 1970. But note that we will compute the cross-sectional standard deviation of firm profits *within* this same balanced panel and compare it to our estimate of common firm-level uncertainty from the panel. Since the two measures are computed over the same panel of firms, any differences between them cannot be attributable to survivorship bias.

Figure 2.10 displays the estimated common uncertainty in firm-level profits $\bar{U}_t^y(h)$ over time for $h = 1, 3$, and 4 quarters. Like the measure of macroeconomic uncertainty analyzed above, these estimates point to a rise in uncertainty surrounding the 1973-75, 1980-82 recessions, but not of the same magnitude. Instead, there are larger increases in common firm-level uncertainty surrounding the 2000-01 and 2007-09 recessions. However, this type of aggregate uncertainty is less countercyclical: the correlation of each of these measures with industrial production growth is negative, but smaller in absolute value than is the correlation of the macro uncertainty measures with production growth. This figure also compares our measures of common firm-level uncertainty $\bar{U}_t^y(h)$ to the popular proxy for common firm-level uncertainty given by on the cross-sectional dispersion in firm profit growth normalized




Figure 2.10: Firm-level Uncertainty: $\bar{U}_t^y(h)$ for $h = 1, 2, 4$. Horizontal lines indicate 1.65 standard deviations above the mean of each series. The thin solid line marked “Dispersion in firm profits” is the cross-sectional standard deviation of firm profit growth, normalized by sales, and denoted \mathcal{D}_t^B . The dispersion is taken after standardizing the profit growth data. Industrial Production (IP) growth is computed as the 12-month moving average of monthly growth rates (in percent). The data are monthly and span the period 1970:Q3-2011:Q2.

by sales, denoted \mathcal{D}_t^B (see equation (2.3)).³⁰

As the figure shows, the two measures behave quite differently, with many more spikes in \mathcal{D}_t^B than in common firm-level uncertainty. Indeed, the dispersion measure exceeds 1.65 standard deviations above its mean dozens of times, while common firm-level uncertainty measures only do so a handful of times. Like the VXO index, there appear to be many movements in the cross-sectional standard deviation of firm profit growth that are not driven by common shocks to uncertainty across firms.

To assess the relative importance of macro uncertainty $\overline{\mathcal{U}}_t^y(h)$ in total uncertainty, we again compute, for each of the 155 firms in the firm-level dataset, and for $h = 1$ to 6, the $R_{jt}^2(h)$ as defined in (2.14), averaged over t . As above, this exercise is performed for the full sample, for recession months, and for non-recession months. Table 2.4 shows that common firm-level uncertainty comprises a larger fraction of the variation in total uncertainty during recessions than during non-recessions, as was the case for common macroeconomic uncertainty. Indeed, the common firm-level common uncertainty we estimate explains an average of 18% of the variation in total uncertainty for an uncertainty horizon of $h = 4$ quarters in non-recessions, but it explains double that in recessions. These results echo those using the macro uncertainty measures. Other results (using VARs for example) are qualitatively similar and omitted to conserve space.

2.7 Conclusion

In this paper we have introduced new time series measures of macroeconomic uncertainty. We have strived to ensure that these measures be comprehensive and as free as possible from both the restrictions of theoretical models and/or dependencies on a handful of economic

³⁰The comparison is affected by how outliers are handled. Bloom (2009) first eliminates all firms with less than 150 quarters of data. He then Winsorizes extreme values that less than the 0.05th percentile and higher than the 99.95th percentile. These percentiles are not taken with respect to the entire distribution of firms. By contrast, we first Winsorize the data and then eliminate firms who have not survived since 1970:Q1. This leaves a considerable number of extreme observations out of our dataset. This has less of an affect on our estimated uncertainty factor, because the data are standardized before estimation.

Average R^2 From regressions of Firm-Level Uncertainty on Common Uncertainty						
Average: $\bar{\mathcal{U}}^y(h) = \frac{1}{N_y} \sum_{j=1}^{N_y} \hat{\mathcal{U}}_{jt}(h)$				PC: $\bar{\mathcal{U}}^y(h) = \sum_{j=1}^{N_y} w_j \hat{\mathcal{U}}_{jt}(h)$		
h	full sample	recession	non-recession	full sample	recession	non-recession
1	0.15	0.29	0.14	0.12	0.27	0.11
2	0.18	0.34	0.16	0.16	0.32	0.14
3	0.19	0.35	0.17	0.17	0.33	0.15
4	0.20	0.36	0.18	0.18	0.33	0.16
5	0.21	0.36	0.18	0.18	0.33	0.16
6	0.21	0.36	0.19	0.18	0.32	0.16

Table 2.4: Values are from regressions of $\hat{\mathcal{U}}_{jt}^y(h)$ on the benchmark (average across series) macro uncertainty measure $\bar{\mathcal{U}}_t^y(h)$ or the principal components (PC) macro uncertainty measure $\bar{\mathcal{U}}_t^y(h)$ over different subsamples. Uncertainty estimated from the quarterly firm-level dataset with observations on firm profit growth rates normalized by sales. Recession months are defined according to the NBER Business Cycle Dating Committee. The data are quarterly and span the period 1970:Q3-2011:Q2.

indicators. We are interested in *macroeconomic* uncertainty, namely uncertainty that may be observed in many economic indicators at the same time, across firms, sectors, markets, and geographic regions. And we are interested in the extent to which this macroeconomic uncertainty is associated with fluctuations in aggregate real activity and financial markets.

Our measures of macroeconomic uncertainty fluctuate in a manner that is often quite distinct from popular proxies for uncertainty, including the volatility of stock market returns (both over time and in the cross-section), the cross-sectional dispersion of firm profits, productivity, or survey-based forecasts. Indeed, our estimates imply far fewer important uncertainty episodes than do popular proxies such as stock market volatility, a measure that forms the basis for the 17 uncertainty dates identified by [Bloom \(2009\)](#). By contrast, we uncover just three big macro uncertainty episodes in the post-war period: the months surrounding the 1973-74 and 1981-82 recessions and the Great Recession of 2007-09, with the 2007-09 recession the most striking episode of heightened uncertainty since 1960. These findings and others reported here suggest that there is much variability in the stock market and in other uncertainty proxies that is not generated by a movement in genuine uncertainty across the broader economy. This occurs both because these proxies over-weight certain

series in the measurement of macro uncertainty, and because they erroneously attribute forecastable fluctuations to a movement in uncertainty.

Our estimates nevertheless point to a quantitatively important dynamic relationship between uncertainty and real activity. In an eleven variable monthly macro VAR, common macro uncertainty shocks have effects on par with monetary policy shocks and are associated with a much larger fraction of the VAR forecast error variance in production and hours worked than are stock market volatility shocks. Our estimates also suggest that macro uncertainty is strongly countercyclical, explaining a much larger component of total uncertainty during recessions than in non-recessions, and far more persistent than common uncertainty proxies.

In this paper we have deliberately taken an atheoretical approach, in order to provide a model-free index of macroeconomic uncertainty that can be tracked over time. Such an index can be used as a benchmark for evaluating any DSGE model with (potentially numerous) primitive stochastic volatility shocks. Our measure of uncertainty conveniently aggregates uncertainty in the economy derived from all sources into one summary statistic. In some cases, it may be useful to construct sub-indices. These can be easily constructed using our framework.

Chapter 3

Volatility, Crises, and Value-at-Risk

3.1 Introduction

Over the past few decades, Value-at-Risk (VaR) has become ubiquitous as a tool for measuring, managing, and regulating the risk exposures of various financial and nonfinancial firms.¹ It started as an esoteric tool employed by quantitative traders to limit the risk on their portfolios during the late 1980s, and made its way into the regulatory environment in 1997 when the U.S. Securities and Exchange Commission (SEC) began requiring firms to disclose information about their market risk, with VaR as one of the accepted alternatives. Since then, VaR has become a centerpiece of the Basel Accords, a series of banking laws and regulations put forth by the Basel Committee on Banking Supervision established in 1974 by the G-10. Despite its widespread adoption, however, VaR has been the subject of much criticism. In response to the proposals leading up to the 2004 Basel II Accord, a handful of academic economists wrote:

“It is our view that the Basel Committee...has failed to address many of the key deficiencies of the global financial regulatory system and even created the

¹Value-at-Risk is defined as the maximum loss on a position or portfolio that can occur over a specified period of time with a small probability. For example, a one-month 99% VaR of \$100 million means that there is only a 1% chance that the value of the portfolio will fall by more than \$ 100 million over the course of the month. For an overview, see [Duffie and Pan \(1997\)](#), or the more recent discussion in [Shin \(2010\)](#).

potential for new sources of instability. The proposed regulations fail to consider that risk is endogenous. Value-at-Risk can destabilize an economy and induce crashes when they would not otherwise occur.” [Dánielsson et al. \(2001\)](#)

In the wake of the recent global financial crisis, these concerns have remained important for market participants, policymakers, and academics alike. In this paper, I shed some new light on the theoretical arguments surrounding this discussion. In particular, I study the equilibrium effects of introducing a VaR constraint into a dynamic economy with homogeneous preferences, endogenous volatility, fire sales, and economically valuable financial intermediation. I find that through its impact on the stationary distribution of wealth in the economy, a VaR constraint can reduce the average level of endogenous volatility, and lower the probability of entering a crisis regime. It does so by forcing agents to sell off their risky asset holdings earlier than they otherwise would, while they have a larger equity buffer to absorb losses. This paper is the first to study the effects of a VaR constraint in a model that does not feature any heterogeneity in preferences or beliefs, and in which endogenous volatility and crises are socially inefficient.

A few existing studies examine the general equilibrium impact of VaR constraints, with somewhat mixed results. They typically all incorporate some form of preference and/or belief heterogeneity, and have nothing to say about whether equilibrium outcomes are socially efficient, either with or without constraints. One exception is [Basak and Shapiro \(2001\)](#), who examine VaR constraints in both partial and general equilibrium settings with homogeneous preferences and find that volatility can endogenously increase when certain agents are VaR-constrained. However, they do not allow the VaR limit to be reevaluated after an initial date, which generates dynamic inconsistencies and stands in contrast to common practice (see [Cuoco et al., 2008](#)). They also do not assume that VaR-constrained agents perform any valuable economic function. [Brunnermeier and Pedersen \(2009\)](#) consider a four-period model with sequentially arriving risk-averse customers, VaR-constrained risk-neutral speculators, and exogenous time-varying fundamental volatility. They find that these features, when

combined with informational imperfections, can lead to situations in which VaR constraints can be destabilizing. [Daníelsson et al. \(2004\)](#) show that VaR constraints can lead to increased price volatility when agents form beliefs based on historical price realizations. In more recent work, [Daníelsson et al. \(2011\)](#) present a model with noise traders and risk-neutral investors which can generate similar conclusions without relying on learning dynamics.

In contrast, [Chabakauri \(2012\)](#) and [Prieto \(2013\)](#) show that in a fully specified dynamic equilibrium model with preference and/or belief heterogeneity, VaR constraints actually reduce endogenous price volatility. The marginal contribution of this paper is to extend these analyses to a setting in which preferences and beliefs are homogeneous, but VaR-constrained agents are assumed to perform a valuable economic function. My results should therefore be understood as complementary to those from existing studies. A goal for future research is to develop a framework where all these different dimensions of heterogeneity can coexist, and provide some quantitative guidance as to which are the most important for the economy as a whole.

The rest of the paper is organized as follows: first, I outline the model and discuss its key features. Second, I introduce a VaR constraint into the model and describe what changes that introduces into the characterization of equilibrium. Third, I simulate the stationary distribution of the model in both cases and examine differences in implications concerning equilibrium volatility. Finally, I perform tests of robustness, discuss some limitations of the analysis, and offer some concluding remarks.

3.2 Unconstrained Model

This section presents the model economy. It is cast in continuous time and features two types of agents: households and experts. These agents are identical in every way, except that experts have access to an economically valuable intermediation technology which allows them to generate higher returns on credit extended to fund productive projects. These

projects are modeled as a single risky asset which represents a claim on the exogenous real dividends they generate.² Both agents are subject to the financial friction that they can only finance these projects by issuing short-term risk-free debt; they cannot issue equity. This can be viewed as an extreme form of the “skin in the game” constraint derived in [He and Krishnamurthy \(2012\)](#), which is necessary to overcome moral hazard on the part of the expert. The value of the project and the risk-free debt security are then determined in equilibrium. As a result of the financial frictions, the equilibrium is a rich one, featuring inefficient countercyclical volatility and destabilizing fire sales.

In terms of machinery, the model falls within the class of continuous-time heterogeneous agent models including [Basak and Cuoco \(1998\)](#), [Longstaff and Wang \(2012\)](#), [Gârleanu and Panageas \(2015\)](#), [Daníelsson et al. \(2011\)](#), [Gârleanu and Pedersen \(2011\)](#), [He and Krishnamurthy \(2013\)](#), and [Brunnermeier and Sannikov \(2014\)](#). The main difference between these models and the one presented here is the coexistence of the following three features: agents do not have heterogeneous beliefs or preferences (either in the form of instantaneous utility or rate of time preference), all agents have access to the same markets (no limited participation is imposed exogenously), and the model gives rise to a stationary distribution which is not degenerate (one type of agent does not dominate the wealth distribution in the long run).

3.2.1 Environment

Time is continuous $t \in [0, \infty)$. There are two types of infinitely-lived agents: households and experts, each of measure one. All agents have identical time-separable preferences over infinite streams of the single perishable consumption good (the numeraire), and discount the future at rate $\rho \in (0, 1)$. They maximize the objectives

$$E \left[\int_0^\infty e^{-\rho t} u(c_t) dt \right], \quad E \left[\int_0^\infty e^{-\rho t} u(\underline{c}_t) dt \right],$$

²I will therefore use “holding the risky asset” and “financing risky projects” interchangeably.

where c_t denotes expert consumption and \underline{c}_t denotes household consumption. I assume that $u(\cdot)$ is strictly increasing, strictly concave, three times continuously differentiable, and satisfies the standard Inada conditions $\lim_{c \rightarrow 0} u'(c) = \infty$, $\lim_{c \rightarrow \infty} u'(c) = 0$. All agents can invest in a risky asset which is in positive net supply, and represents claims to the exogenous real dividends generated by productive project (as in [Dánielsson et al., 2011](#)). While both households and experts can fund projects, households are less effective at doing so. Experts have access to an economically valuable intermediation technology which enables project managers to make better use of credit.³ When projects are funded by experts, they generate a real dividend $a\delta_t$ per unit of time, where $a > 0$ is a scalar, and δ_t follows

$$\frac{d\delta_t}{\delta_t} = gdt + \sigma dZ_t.$$

Z_t is a standard Brownian motion on a complete probability space, and represents the only source of uncertainty in the economy; g and σ are scalars. On the other hand, when projects are funded through credit directly extended by households, they generate a real dividend $\underline{a}\delta_t$ where $\underline{a} < a$, and $\underline{\delta}_t$ follows

$$\frac{d\underline{\delta}_t}{\underline{\delta}_t} = \underline{g}dt + \sigma dZ_t,$$

with $\underline{g} < g$. This difference in efficiency is the only source of heterogeneity in the model. It is important at this point to notice that in existing general equilibrium analyses of risk-based constraints, there is essentially no economically valuable function performed by the constrained agent. For example, [Dánielsson et al. \(2011\)](#), [Prieto \(2013\)](#), and [Chabakauri \(2012\)](#) all assume that the only difference between households and intermediaries lies in preferences (or beliefs). While it is certainly plausible that attitudes towards risk are important dimensions of agent heterogeneity, it also seems important to remember that intermediaries

³The virtue of this specification is its generality. Many different stories can be told for why intermediation is valuable, and extending the model to more carefully incorporate various mechanisms would be a fruitful extension. For example: (i) funding risky projects requires some degree of financial “sophistication.” When end-use borrowers are forced to obtain funding from unsophisticated households, resources are wasted, (ii) intermediaries have access to a better monitoring technology which can be used to extract higher returns or exert influence on project management.

do perform an economically valuable function in the economy. This is especially worth emphasizing in light of the recent crisis, where the inability of intermediaries to perform this function took center stage. The structure assumed here is meant to capture the value of intermediation in a reduced-form way.

Letting q_t denote the price of the risky asset in terms of the consumption good, the total returns on the risky asset for the expert and household, respectively, are defined as:

$$dr_t^D = \frac{a\delta_t dt + d(q_t\delta_t)}{q_t\delta_t}, \quad dr_t^{\underline{D}} = \frac{\underline{a}\delta_t dt + d(q_t\underline{\delta}_t)}{q_t\underline{\delta}_t}.$$

Agents also have access to a risk-free bond which is in zero net supply and generates a return denoted by dr_t . Neither households nor experts can issue equity. The only way they can finance asset holdings greater than their net worth is by shorting the risk-free asset (borrowing through risk-free debt). Since experts will be net borrowers in equilibrium, resources invested in this asset can represent (risk-free) deposits made by the household with the expert-owned intermediary. However, agents are restricted from shorting the risky asset, which is natural given its interpretation as a risky project.⁴

Finally, experts and households each exogenously receive labor income of θD_t per unit of time, where $\theta > 0$ is a scalar and D_t is the aggregate dividend on the risky asset at time t . As in [He and Krishnamurthy \(2013\)](#), I assume that labor income is proportional to dividends to preserve the scale-invariance of the economy. The assumption of exogenous labor income is crucial for obtaining a stationary distribution that is not degenerate. Without it, experts dominate the economy and households disappear in the long run (cf. [Dumas, 1989](#); [Wang, 1996](#)). However, by assuming that the income flow arrives at the same rate for both agents, I ensure that this does not introduce another dimension of heterogeneity.⁵ Lastly, I assume that agents ignore this exogenous income process when solving their optimization programs.

⁴This no short sale constraint is not a necessary feature of the model, but it does help facilitate the interpretation of results.

⁵Of course, I could allow for heterogeneity here by introducing two parameters, θ and $\underline{\theta}$, instead of one. I allow for this possibility later on in a robustness check.

This assumption guarantees that the optimal policy rules of households and experts can still be expressed in closed form.⁶ Given the two assets and exogenous labor income, expert and household net worth, respectively, evolve according to the following diffusion processes in equilibrium:

$$\begin{aligned} dN_t &= \underbrace{q_t \delta_t dr_t^D}_{\text{risky investment}} - \underbrace{(q_t \delta_t - N_t) dr_t}_{\text{risk-free debt}} - \underbrace{c_t dt}_{\text{consumption}} + \underbrace{\theta D_t dt}_{\text{labor income}}, \\ d\underline{N}_t &= \underbrace{q_t \underline{\delta}_t d\underline{r}_t^D}_{\text{risky investment}} + \underbrace{(\underline{N}_t - q_t \underline{\delta}_t) dr_t}_{\text{risk-free savings}} - \underbrace{\underline{c}_t dt}_{\text{consumption}} + \underbrace{\theta D_t dt}_{\text{labor income}}. \end{aligned}$$

Given the structural environment outlined above, equilibrium is defined in the following way:

Definition 2. *An equilibrium is a set of price processes $\{dq_t\}$ and $\{dr_t\}$, and allocations $\{c_t, \underline{c}_t, \delta_t, \underline{\delta}_t\}$, such that given initial allocations and prices:⁷*

1. Expert optimality: $\{c_t, \delta_t\}$ solves

$$\max_{c_t, \delta_t \geq 0} E \left[\int_0^\infty e^{-\rho t} u(c_t) dt \right] \quad s.t.$$

$$\begin{aligned} dN_t &= a \delta_t dt + d(q_t \delta_t) - (q_t \delta_t - N_t) dr_t - c_t dt \\ N_t &\geq 0. \end{aligned}$$

2. Household optimality: $\{\underline{c}_t, \underline{\delta}_t\}$ solves

$$\max_{\underline{c}_t, \underline{\delta}_t \geq 0} E \left[\int_0^\infty e^{-\rho t} u(\underline{c}_t) dt \right] \quad s.t.$$

⁶Alternatively, it is possible to model agents using an overlapping generations structure, as in [He and Krishnamurthy \(2013\)](#). In their model, a unit mass of generation t agents are born with wealth N_t and live from period t to period $t + \Delta$. They maximize utility $\rho \Delta u(c_t) + (1 - \rho \Delta) E_t [b(N_{t+\Delta})]$, where $u(\cdot)$ is a felicity function and $b(\cdot)$ is a bequest function. Choosing $u(x) = b(x) = \log(x)$ and letting $\Delta \rightarrow dt$ delivers exactly the same results as the ones reported in the following sections.

⁷Technical measurability and integrability conditions of $\{c_t, \underline{c}_t, \delta_t, \underline{\delta}_t\}$ also must hold. Furthermore, notice that the exogenous income term θD_t does not appear in the agents' programs, but is required to appear in the equilibrium laws of motion for agent net worth.

$$\begin{aligned}
d\underline{N}_t &= \underline{a}\delta_t dt + d(q_t \underline{\delta}_t) - (q_t \underline{\delta}_t - \underline{N}_t)dr_t - \underline{c}_t dt \\
\underline{N}_t &\geq 0.
\end{aligned}$$

3. Market clearing: the risky asset and goods market clear⁸

$$\begin{aligned}
\delta_t + \underline{\delta}_t &= D_t, \\
c_t + \underline{c}_t &= a\delta_t + \underline{a}\delta_t + 2\theta D_t.
\end{aligned}$$

4. Laws of motion: the aggregate project dividend, expert net worth, and household net worth, respectively, follow

$$\begin{aligned}
dD_t &= (g\delta_t + \underline{g}\delta_t)dt + \sigma D_t dZ_t \\
dN_t &= a\delta_t dt + d(q_t \delta_t) - (q_t \delta_t - N_t)dr_t - c_t dt + \theta D_t dt \\
d\underline{N}_t &= \underline{a}\delta_t dt + d(q_t \underline{\delta}_t) - (q_t \underline{\delta}_t - \underline{N}_t)dr_t - \underline{c}_t dt + \theta D_t dt.
\end{aligned}$$

3.2.2 Pareto optimality

Before solving for the competitive outcome of the above economy, I first present the first-best outcome, which will serve as a useful benchmark for understanding the various externalities and inefficiencies that arise in the decentralized version. The planner's problem in this case would be to maximize a weighted sum of expert and household utilities, subject to the two intermediation technologies, the resource constraint, and the short-sale constraint.⁹ This implies that the only advantage that the planner has is to be able to freely transfer resources across agents without being subject to the no-equity financial constraint. The solution to this problem takes a particularly simple form: intuitively, it is socially advantageous for

⁸Since each type actually represents a continuum of agents, δ_t and $\underline{\delta}_t$ represent risky asset holdings integrated over the entire measure of agents. The same is true for c_t and \underline{c}_t . To avoid notational clutter, I leave this integration implicit throughout. Also note that the risk-free asset market clears by Walras's Law.

⁹This is necessary so that the planner's problem has a well-defined solution.

all credit to be intermediated, since households are less efficient at channeling funds to the ultimate borrowers. The resulting project payoffs (and exogenous labor income) can then be then distributed to the two types of agents according to their respective weights in the planner's objective. Formally, let $\tau \in [0, 1]$ and $1 - \tau$ denote the ex-ante Pareto weights attached to the expert and household sectors, respectively. Then, the planner's problem is

$$\max_{c_t, \underline{c}_t, \delta_t, \underline{\delta}_t \geq 0} E \left[\int_0^\infty e^{-\rho t} \{ \tau u(c_t) + (1 - \tau) u(\underline{c}_t) \} dt \right], \quad s.t.$$

$$\begin{aligned} dD_t &= (g\delta_t + \underline{g}\underline{\delta}_t)dt + \sigma D_t dZ_t \\ \delta_t + \underline{\delta}_t &= D_t \\ c_t + \underline{c}_t &= a\delta_t + \underline{a}\underline{\delta}_t + 2\theta D_t. \end{aligned}$$

The following proposition demonstrates that the intuition outlined above is indeed correct.¹⁰

Proposition 1. *The solution to the planner's problem is*

$$\begin{aligned} c_t &= \tau(a + 2\theta)D_t & \underline{c}_t &= (1 - \tau)(a + 2\theta)D_t \\ \delta_t &= D_t & \underline{\delta}_t &= 0. \end{aligned}$$

3.2.3 Competitive equilibrium

In this section, I show how the equilibrium of the unconstrained economy can be fully characterized by a system of nonlinear ordinary differential equations (ODE). Due to the form of financial constraints and agent heterogeneity in the model, the aggregation results developed by [Cuoco and He \(2001\)](#) and employed in [Basak and Cuoco \(1998\)](#) are not readily applicable. This is because the consumption of the weakly-aggregated representative agent is not equal to the (exogenous) aggregate dividend D_t due to differences in intermediation

¹⁰All proofs are included in Appendix [C](#).

technology. The solution method therefore closely follows the one used in [Brunnermeier and Sannikov \(2014\)](#). In particular, I will look for a Markov equilibrium in which the price of the risky asset follows

$$\frac{dq_t}{q_t} = \mu_t^q dt + \sigma_t^q dZ_t,$$

where μ_t^q σ_t^q are determined in equilibrium. Moreover, to greatly facilitate the following exposition, I will assume that $u(c) = \log(c)$. The form of this function could be straightforwardly generalized, for example, to the CRRA class $u(c) = c^{1-\gamma}/(1-\gamma)$ where $\gamma > 0$ without substantially altering the results. What matters is just that γ is the same for both agents.

Agent optimality

Using the law of motion for q_t , Itô's Lemma, and the definition above, the return on the risky asset for experts can be written as

$$dr_t^D = \underbrace{a/q_t dt}_{\text{dividend yield}} + \underbrace{(g + \mu_t^q + \sigma \sigma_t^q) dt + (\sigma + \sigma_t^q) dZ_t}_{\text{capital gains}}.$$

The important feature to notice at this point is that the risk of this asset can be cleanly decomposed into two parts: a fundamental component, σ , which purely depends on the law of motion of the exogenous dividend process, and an endogenous component, σ_t^q , which depends on the equilibrium behavior of agents in the economy. An interesting feature of this model is that as a result of financial frictions, σ_t^q will not in general be zero, but will fluctuate countercyclically (and asymmetrically) over time, in line with the empirical evidence.¹¹ These random price movements directly affect the net worth of experts, which

¹¹[Mele \(2007\)](#) studies the properties necessary for σ_t^q to exhibit this behavior in a similar class of economies. He also documents that such behavior is broadly consistent with the historical fluctuations of the price-dividend ratio on the Standard & Poor's (S&P) Composite index.

can be seen by substituting the above expression for dr_t^D into the law of motion of net worth:

$$\begin{aligned} dN_t = & q_t \delta_t (a/q_t + g + \mu_t^q + \sigma \sigma_t^q) - (q_t \delta_t - N_t) dr_t \\ & - (c_t + \theta D_t) dt + q_t \delta_t (\sigma + \sigma_t^q) dZ_t. \end{aligned}$$

Similarly, the return on the risky asset for households is

$$d\bar{r}_t^D = (\bar{a}/q_t + \bar{g} + \mu_t^q + \sigma \sigma_t^q) dt + (\sigma + \sigma_t^q) dZ_t.$$

Since both households and experts have the same preferences, access to the same securities, and face the same constraints on their consumption and portfolio choice, their optimal policies are also the same.

Proposition 2. *The solutions to the household and expert problems are given by the time-invariant policy rules:*

$$\begin{aligned} c_t &= \rho N_t, & \underline{c}_t &= \rho \underline{N}_t, \\ q_t \delta_t (\sigma + \sigma_t^q)^2 dt &\geq E_t(dr_t^D - dr_t) N_t, & q_t \delta_t (\sigma + \sigma_t^q)^2 dt &\geq E_t(d\bar{r}_t^D - dr_t) \underline{N}_t, \end{aligned}$$

where the second condition holds with equality for experts whenever $\delta_t > 0$, and for households whenever $\delta_t > 0$.

This proposition states that experts consume a constant fraction of their net worth, which is a convenient result implied by logarithmic preferences (cf. [Merton, 1969, 1971](#)), and invest only up until the point that their expected excess return on the risky asset is equal to a rescaling of the riskiness of their net worth.

State evolution

Since agents are heterogeneous and face financial frictions, the distribution of wealth matters for the determination of equilibrium prices and allocations. This insight is well-known in

the macroeconomic literature, going back at least to [Krusell and Smith \(1998\)](#), who tackled this problem in a discrete-time RBC setting. In the context of this model, the distribution of wealth is relatively easy to keep track of, since there are only two types of agents. As a summary of this two-point distribution, define the state variable

$$\eta_t \equiv \frac{N_t}{q_t D_t} \in [0, 1],$$

which is the fraction of wealth held by experts at time t . Note also that since bonds are in zero net supply, the net worth of experts and households must sum to the value of the risky asset: $N_t + \underline{N}_t = q_t D_t$, so that the fraction of wealth held by households is simply $1 - \eta_t$. Also, it will be convenient to let $\psi_t \equiv \delta_t / D_t$ denote the fraction of the risky asset intermediated by experts, and $1 - \psi_t$ the fraction financed directly by households. Using Itô's lemma, this state variable has law of motion:

$$d\eta_t = dN_t \frac{1}{q_t D_t} + d \left(\frac{1}{q_t D_t} \right) N_t - dN_t d \left(\frac{1}{q_t D_t} \right),$$

where

$$\begin{aligned} dN_t &= \psi_t q_t D_t (\underline{a}/q_t + \underline{g} + \mu_t^q + \sigma \sigma_t^q) dt - \rho N_t dt + \theta D_t dt \\ &\quad - (\psi_t q_t D_t - N_t) dr_t + \psi_t q_t D_t (\sigma + \sigma_t^q) dZ_t, \\ d \left(\frac{1/(q_t D_t)}{1/(q_t D_t)} \right) &= -(g\psi_t + \underline{g}(1 - \psi_t) + \mu_t^q + \sigma_t^q) dt - (\sigma + \sigma_t^q) dZ_t. \end{aligned}$$

Combining expressions implies that

$$\begin{aligned} d\eta_t &= [(\psi_t - \eta_t)(\mu_t^q + \sigma \sigma_t^q - (\sigma + \sigma_t^q)^2) + \psi_t(\underline{a}/q_t + \underline{g}) + \theta/q_t \\ &\quad - (\rho + g\psi_t + \underline{g}(1 - \psi_t))\eta_t] dt - (\psi_t - \eta_t) dr_t + (\psi_t - \eta_t)(\sigma + \sigma_t^q) dZ_t. \end{aligned}$$

This says that the distribution of wealth depends on fundamental shocks as long as $\psi_t \neq \eta_t$. Since in general, experts will want to be leveraged, this condition is always satisfied in the interior of the state space. Define the diffusion term $\sigma_t^\eta \equiv (\psi_t - \eta_t)(\sigma + \sigma_t^q)$ so that $d\eta_t = E_t(d\eta_t) + \sigma_t^\eta dZ_t$.

Market clearing and prices

Given the law of motion for the state variable η_t it is necessary to solve for the equilibrium price function $q_t = q(\eta_t)$. To do this, first note that since $N_t + \underline{N}_t = q_t D_t$, it follows from agent optimality that aggregate consumption is given by $c_t + \underline{c}_t = \rho q_t D_t$. Combining this with market clearing in the goods market implies that $\rho q_t D_t = (a\psi_t + \underline{a}(1 - \psi_t))D_t + 2\theta D_t$, and therefore that

$$q_t = \frac{\psi_t a + (1 - \psi_t)\underline{a} + 2\theta}{\rho}.$$

This key equilibrium relationship implies that the price of the risky asset is increasing in the fraction of the asset held by experts (since $a > \underline{a}$). Intuitively, the price reflects the marginal value of the asset in the economy. If there is an equilibrium outcome where it is optimal for experts finance all of the risky projects, $\psi_t = 1$, it must be that the price of the asset only reflects the marginal valuation of experts, adjusted to account for non-asset income: $q_t = \bar{q} \equiv (a + 2\theta)/\rho$. On the other hand, if it is optimal for households to finance all of the risky projects directly, the price must adjust to ensure that $q_t = \underline{q} \equiv (\underline{a} + 2\theta)/\rho < \bar{q}$. In general, the price will fluctuate over the interval $[\underline{q}, \bar{q}]$ depending on which agent is carrying out more of the financing. The important point here is that this is an equilibrium relationship which is *not internalized* by agents while they solve their intertemporal optimization problems. As I will show below, it turns out that the failure of agents to internalize this pecuniary effect generates an externality where inefficient, destabilizing fire sales can occur. For now, rewrite the above equation to express $\psi_t = \psi(\eta_t)$ as a function of $q_t = q(\eta_t)$,

$$\psi(\eta_t) = \frac{\rho q(\eta_t) - \underline{a} - 2\theta}{a - \underline{a}}.$$

Next, notice that by applying Itô's lemma to the function $q(\eta_t)$ and matching coefficients with the law of motion of q_t , it must be that

$$\begin{aligned}\sigma_t^q q(\eta_t) &= \sigma_t^\eta q'(\eta_t) \\ &= (\psi(\eta_t) - \eta_t)(\sigma + \sigma_t^q)q'(\eta_t),\end{aligned}$$

where the second equality comes from the law of motion of $d\eta_t$ derived above. This expression represents a nonlinear first-order ordinary differential equation (ODE) which fully characterizes q_t as a function of the state variable η_t , as long as it is possible to find an expression for the function $\sigma_t^q = \sigma^q(\eta_t)$. The following proposition shows that it is indeed possible to do this, and describes the evolution of all prices as a function of η_t .

Proposition 3. *Whenever the household finds it optimal to finance some of the risky projects in equilibrium, $(1 - \psi_t > 0)$, the laws of motion for prices and the state variable are:*

$$\begin{aligned}dr_t &= r_t dt \\ dq_t &= \mu_t^q q_t dt + \sigma^q q_t dZ_t \\ d\eta_t &= \mu_t^\eta dt + \sigma_t^\eta dZ_t,\end{aligned}$$

where the functions $\sigma_t^q = \sigma^q(\eta_t)$, $\mu_t^q = \mu^q(\eta_t)$, $\mu_t^\eta = \mu^\eta(\eta_t)$, $\sigma_t^\eta = \sigma^\eta(\eta_t)$, and $r_t = r(\eta_t)$ satisfy the system of (deterministic) nonlinear ordinary differential equations:

$$\begin{aligned}\psi(\eta) &= \frac{\rho q(\eta) - \underline{a} - 2\theta}{a - \underline{a}} \\ \sigma^q(\eta) &= \left[\frac{(a - \underline{a})/q(\eta) + (g - \underline{g})}{\psi(\eta)/\eta - (1 - \psi(\eta))/(1 - \eta)} \right]^{1/2} - \sigma \\ q'(\eta) &= \left(1 - \frac{\sigma}{\sigma + \sigma^q(\eta)} \right) \frac{q(\eta)}{\psi(\eta) - \eta}\end{aligned}$$

$$\begin{aligned}
\mu^\eta(\eta) &= \psi(\eta)(\underline{a}/q(\eta) + \underline{g}) - (\psi(\eta) - \eta)(\underline{a}/q(\eta) + \underline{g}) - [\rho + g\psi(\eta) \\
&\quad + \underline{g}(1 - \psi(\eta))]\eta + \theta/q(\eta) - \frac{(\psi(\eta) - \eta)^2}{(1 - \eta)}(\sigma + \sigma^q(\eta))^2 \\
\sigma^\eta(\eta) &= (\psi(\eta) - \eta)(\sigma + \sigma^q(\eta)) \\
\mu^q(\eta) &= \mu^\eta(\eta)q'(\eta) + \frac{1}{2}(\psi(\eta) - \eta)^2(\sigma + \sigma^q(\eta))^2q''(\eta) \\
r(\eta) &= \mu^q(\eta) + \sigma\sigma^q(\eta) + \underline{a}/q(\eta) + \underline{g} - \left(\frac{1 - \psi(\eta)}{1 - \eta}\right)(\sigma + \sigma^q(\eta))^2
\end{aligned}$$

Whenever the household does not find it optimal to finance some of the risky projects ($\psi_t = 1$), the equilibrium price is $q_t = \bar{q} = (a + 2\theta)/\rho$, so $\mu_t^q = \sigma_t^q = 0$, and the laws of motion for the interest rate and state variable are given by.¹²

$$\begin{aligned}
dr_t &= r_t dt \\
d\eta_t &= \mu_t^\eta dt + \sigma_t^\eta dZ_t,
\end{aligned}$$

where the functions $\mu_t^\eta = \mu^\eta(\eta_t)$, $\sigma_t^\eta = \sigma^\eta(\eta_t)$, and $r_t = r(\eta_t)$ have the closed-form expressions:

$$\begin{aligned}
\mu^\eta(\eta) &= \frac{(1 - \eta)^2}{\eta}\sigma^2 + \rho \left[\left(\frac{a}{a + 2\theta} - 1 \right) \eta_t + \frac{\theta}{a + 2\theta} \right] \\
\sigma^\eta(\eta) &= (1 - \eta)\sigma \\
r(\eta) &= \frac{a\rho}{a + 2\theta} + g - \frac{\sigma^2}{\eta}.
\end{aligned}$$

Furthermore, defining $\eta^* \equiv \sup\{\eta \in [0, 1] \mid \psi(\eta^*) < 1\}$, the equilibrium is also subject to the following three boundary conditions

$$q(0) = \frac{(\underline{a} + 2\theta)}{\rho}, \quad q(1) = \frac{(a + 2\theta)}{\rho}, \quad q(\eta^*) = q(1).$$

¹²When short sales are allowed, it will never be the case that $\psi_t = 1$ in equilibrium. Therefore only the first part of this proposition is relevant in that case, and the boundary condition $q(\eta^*) = q(1)$ is redundant.

3.2.4 Numerical example

The equilibrium can be computed numerically using an ODE solver such as MATLAB's `ode45.m`.¹³ The following numerical example illustrates some of the equilibrium properties of the model. In this example, I select the following parameter values: $a = 1$, $\underline{a} = 0.25$, $\rho = 0.05$, $\sigma = 0.06$, $g = 0.05$, $\underline{g} = 0.01$, $\theta = 0.05$. In Section 3.4, I solve the model many times for a set of different parameter constellations for robustness. Figure 3.1 plots the policy functions $q(\eta)$, $\psi(\eta)$, $\mu^\eta(\eta)$, and $\sigma^\eta(\eta)$.

The economy features two regimes: one in which experts are well capitalized (η_t is high) and intermediate all financing in the economy. Prices are stable and the economy has relatively low volatility, and does not change rapidly on average. The other is one in which experts are poorly capitalized (η_t is low), and are unwilling to extend sufficient credit to finance risky projects, so households step in and provide (inefficient) direct credit. As a result, there is a fire sale of assets, a decline in prices, and a sharp increase in the volatility and expected rate of change of the economy. Because this strongly resonates with what we observe in practice during crisis episodes, I define the (endogenous) set of crisis states as

$$\Theta(\eta) \equiv \{\eta \in [0, 1] \mid \eta < \eta^*\}$$

where η^* is the point of regime shift. In a later section, I will be interested in determining $P(\eta \in \Theta(\eta))$, the probability that the economy finds itself in a crisis, and how this object changes once VaR constraints are introduced. Note that $P(\cdot)$ denotes the probability taken with respect to the stationary distribution implied by equilibrium.

The driving force behind these various nonlinear price effects is the coalescence of heterogeneous intermediation technologies and financial constraints. The value of assets on the balance sheets of experts depends on the price of capital, but liabilities do not, since all debt is risk-free and none of the agents are permitted to issue equity claims. When the net

¹³The MATLAB code used to solve this model are built off of the suite of programs provided by [Sannikov \(2013\)](#).

worth of experts is low, they may optimally choose not to extend so much risky credit, since they are risk averse. However, they do not take into account that whenever they choose not to completely finance the risky project, households will be forced to do so (by market clearing), which will lead to a fall in the price of risky loans. This price drop adversely impacts experts because it means they are subject to further losses on their assets. An adverse feedback loop ensues generating inefficiently excessive volatility, and wasted real resources. These dynamics are also present in the model of [Brunnermeier and Sannikov \(2014\)](#), who provide a more detailed discussion of this amplification mechanism; they reference [Lorenzoni \(2008\)](#) and [Jeanne and Korinek \(2010\)](#) as other recent applications of this type of pecuniary externality. One important difference between this equilibrium and theirs is the behavior of the state drift, μ_t^η , which they find to be nonnegative over the entire state space, but which can be negative in this case. The reason for this involves my assumptions about labor income, which I discuss in [Section 3.4](#).

3.3 Adding Value-at Risk

This section modifies the above economy by superimposing a Value-at-Risk (VaR) constraint on each individual expert. Formally, the constraint requires that each expert hold a fraction $\beta \geq 0$ of his net worth in order to cover his π_v -probability VaR, v_t . To better understand the VaR constraint, index time as $t, t + \Delta t, t + 2\Delta t, \dots$. Define $\Delta N_t \equiv (N_{t+\Delta t} - N_t)/\sqrt{\Delta t}$ to be the normalized change in expert net worth from t to $t + \Delta t$. The VaR constraint requires that $v_t \leq \beta N_t$, where v_t is defined as

$$v_t \equiv \inf\{v \geq 0 \mid P(\Delta N_t \leq -v) = \pi_v\}.$$

This says that v_t is the minimum amount of equity needed to ensure that losses greater than v_t per $\sqrt{\Delta t}$ units of time only happen with a small probability π_v . The next proposition shows that this constraint takes a simple form in the model of the previous section (the same

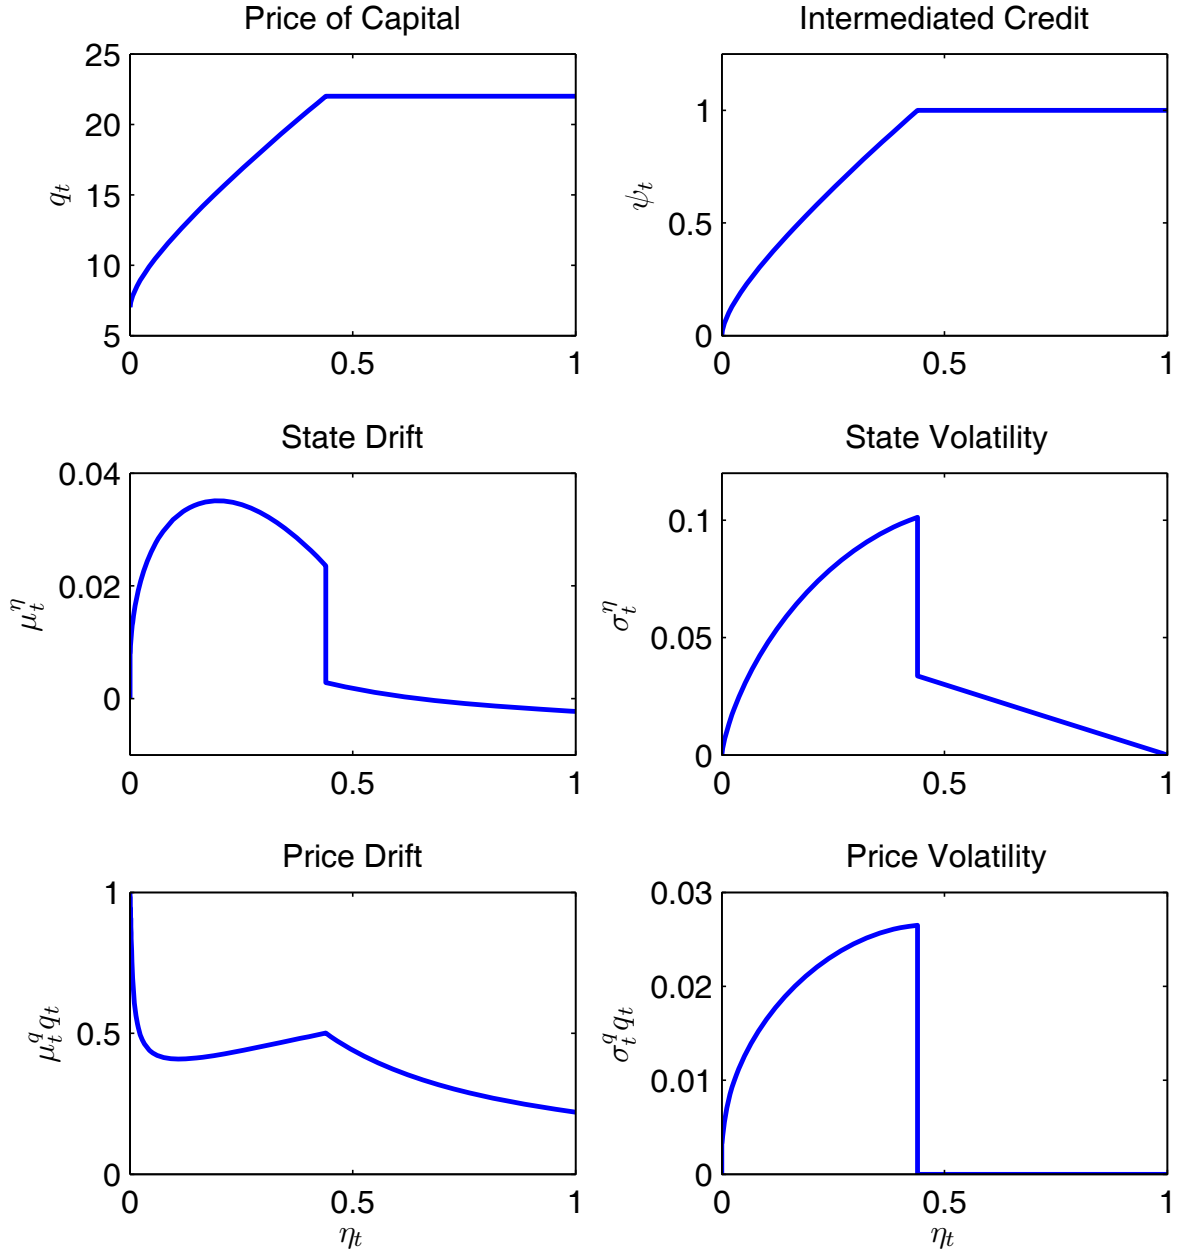


Figure 3.1: Equilibrium variables as functions of η_t in the unconstrained equilibrium. The economy features two regimes: “normal times,” when experts are well capitalized and volatility is low, and “crisis episodes,” when experts are poorly capitalized, and volatility is high.

form used by [Gârleanu and Pedersen \(2011\)](#), [Prieto \(2013\)](#), and [Dánielsson et al. \(2011\)](#).

Proposition 4. *Let v_t denote an individual expert's π_v -probability VaR, and define $\alpha \equiv \Phi^{-1}(1 - \pi_v)/\beta \geq 0$, where $\Phi(\cdot)$ is the standard normal cumulative distribution function. As $\Delta t \rightarrow 0$*

$$v_t \leq \beta N_t \Leftrightarrow \alpha q_t \delta_t (\sigma + \sigma_t^q) \leq N_t.$$

Therefore, the VaR constraint restricts the leverage of each expert to be no larger than $1/(\alpha(\sigma + \sigma_t^q))$. Because all debt is risk-free in this model, the left hand side of the constraint is a multiple of the total volatility of the expert's asset holdings. Therefore, the expert must set aside more net worth as the volatility of his asset holdings increases. In what follows, I impose this constraint on the expert, and solve for the equilibrium as above.

3.3.1 Constrained equilibrium

The new problem for experts is to solve

$$\max_{c_t, \delta_t \geq 0} E \left[\int_0^\infty e^{-\rho t} u(c_t) dt \right] \quad s.t.$$

$$dN_t = a\delta_t dt + d(q_t \delta_t) - (q_t \delta_t - N_t)dr_t - c_t dt$$

$$N_t \geq \alpha q_t \delta_t (\sigma + \sigma_t^q).$$

The rest of the economy is unchanged, and the resulting equilibrium can be defined and computed exactly as above. The only thing needed is to find the new optimal consumption and portfolio allocation rules of the expert under the occasionally binding VaR constraint. As it turns out, the form of these rules can still be obtained in closed form under CRRA preferences of the form $u(c) = c^{1-\gamma}/(1-\gamma)$ with $\gamma > 0$. For simplicity, assume as above that

$\gamma \rightarrow 1$ so that $u(c) = \log(c)$. In the case of logarithmic utility, the expert portfolio decision is equivalent to solving

$$\begin{aligned} \max_{\delta_t \geq 0} & \left\{ q_t \delta_t E_t(dr_t^D - dr_t) N_t - \frac{1}{2} (q_t \delta_t)^2 \text{var}_t(dr_t^D) \right\} \\ \text{s.t.} & \quad \alpha q_t \delta_t (\sigma + \sigma_t^q) \leq N_t. \end{aligned}$$

The optimality conditions associated with this problem are similar to the unconstrained case, except now an additional Lagrange multiplier enters, and a complementary slackness condition is included. Let $\lambda_t/(\alpha(\sigma + \sigma_t^q))$ denote the Lagrange multiplier on the constraint. Then

$$\begin{aligned} q_t \delta_t (\sigma + \sigma_t^q)^2 &= E_t(dr_t^D - dr_t) N_t - \lambda_t dt, \\ \lambda_t (N_t - \alpha q_t \delta_t (\sigma + \sigma_t^q)) &= 0. \end{aligned}$$

Defining $A_t \equiv 1 + \lambda_t N_t / [q_t \delta_t (\sigma + \sigma_t^q)^2]$, the optimality condition can be rewritten as:

$$q_t \delta_t (\sigma + \sigma_t^q)^2 A_t \leq E_t(dr_t^k - dr_t) N_t.$$

This looks like the standard portfolio allocation solution between one risky and one risk free asset (since $q_t \delta_t / N_t$ denotes the fraction of wealth invested in the risky asset). In this case, however, the volatility constraint introduces what looks like time variation in risk aversion, A_t . This point has been made extensively in a static context by [Shin \(2010\)](#), and in a dynamic context by [Dánielsson et al. \(2011\)](#). The following result summarizes the expert solution.

Proposition 5. *The solution to the expert problem under a VaR constraint is given by the*

time-invariant policy rules:

$$\begin{aligned} c_t &= \rho N_t \\ q_t \delta_t (\sigma + \sigma_t^q)^2 &\leq \min \left\{ E_t(dr_t^D - dr_t) N_t, \frac{1}{\alpha(\sigma + \sigma_t^q)} N_t \right\} \end{aligned}$$

where the second condition holds with equality whenever $\delta_t > 0$.

Since the goal of the paper is to isolate the effect of this VaR constraint, I leave everything else exactly as in the unconstrained case. For completeness, the next proposition is a modification of Proposition 3 for the constrained economy. In particular, note that in states of the world where it is optimal for experts to exclusively handle the intermediation in the economy (i.e. $\psi_t = 1$) it implies that the VaR constraint is not binding. This result has an intuitive appeal, since it means that VaR constraints are not binding during normal times.

Proposition 6. *Whenever the household finds it optimal to finance some of the risky projects in equilibrium, ($1 - \psi_t > 0$), the laws of motion for prices and the state variable are:*

$$\begin{aligned} dr_t &= r_t dt \\ dq_t &= \mu_t^q q_t dt + \sigma^q q_t dZ_t \\ d\eta_t &= \mu_t^\eta dt + \sigma_t^\eta dZ_t, \end{aligned}$$

where the functions $\sigma_t^q = \sigma^q(\eta_t)$, $\mu_t^q = \mu^q(\eta_t)$, $\mu_t^\eta = \mu^\eta(\eta_t)$, $\sigma_t^\eta = \sigma^\eta(\eta_t)$, and $r_t = r(\eta_t)$ satisfy the system of (deterministic) nonlinear ordinary differential equations:

$$\begin{aligned} \psi(\eta) &= \frac{\rho q(\eta) - \underline{a} - 2\theta}{a - \underline{a}} \\ \psi(\eta) &= \min \left\{ \frac{[(a - \underline{a})/q(\eta) + (g - \underline{g})]\eta}{(\sigma + \sigma^q(\eta))^2} + \frac{1 - \psi(\eta)}{1 - \eta}, \frac{1}{\alpha(\sigma + \sigma^q(\eta))} \right\} \\ q'(\eta) &= \left(1 - \frac{\sigma}{\sigma + \sigma^q(\eta)} \right) \frac{q(\eta)}{\psi(\eta) - \eta} \end{aligned}$$

$$\begin{aligned}
\mu^q(\eta) &= \mu^\eta(\eta)q'(\eta) + \frac{1}{2}(\psi(\eta) - \eta)^2(\sigma + \sigma^q(\eta))^2q''(\eta), \\
\mu^\eta(\eta) &= \psi(\eta)(\underline{a}/q(\eta) + g) - (\psi(\eta) - \eta)(\underline{a}/q(\eta) + \underline{g}) - [\rho + g\psi(\eta) \\
&\quad + \underline{g}(1 - \psi(\eta))]\eta + \theta/q(\eta) - \frac{(\psi(\eta) - \eta)^2}{(1 - \eta)}(\sigma + \sigma^q(\eta))^2 \\
\sigma^\eta(\eta) &= (\psi(\eta) - \eta)(\sigma + \sigma^q(\eta)) \\
r(\eta) &= \mu^q(\eta) + \sigma\sigma^q(\eta) + \underline{a}/q(\eta) + \underline{g} - \left(\frac{1 - \psi(\eta)}{1 - \eta}\right)(\sigma + \sigma^q(\eta))^2
\end{aligned}$$

Whenever the household does not find it optimal to finance some of the risky projects ($\psi_t = 1$), and as long as $\alpha\sigma \neq 1$, the equilibrium price is $q_t = \bar{q} = (a + 2\theta)/\rho$, so $\mu_t^q = \sigma_t^q = 0$, and the laws of motion for the interest rate and state variable are given by:

$$\begin{aligned}
dr_t &= r_t dt \\
d\eta_t &= \mu_t^\eta dt + \sigma_t^\eta dZ_t,
\end{aligned}$$

where the functions $\mu_t^\eta = \mu^\eta(\eta_t)$, $\sigma_t^\eta = \sigma^\eta(\eta_t)$, and $r_t = r(\eta_t)$ have the closed-form expressions:

$$\begin{aligned}
\mu^\eta(\eta) &= \frac{(1 - \eta)^2}{\eta}\sigma^2 + \rho \left[\left(\frac{a}{a + 2\theta} - 1 \right) \eta + \frac{\theta}{a + 2\theta} \right] \\
\sigma^\eta(\eta) &= (1 - \eta)\sigma \\
r(\eta) &= \frac{a\rho}{a + 2\theta} + g - \frac{\sigma^2}{\eta}.
\end{aligned}$$

Furthermore, defining $\eta^* \equiv \sup\{\eta \in [0, 1] \mid \psi(\eta^*) < 1\}$, the equilibrium is also subject to the following three boundary conditions

$$q(0) = \frac{(\underline{a} + 2\theta)}{\rho}, \quad q(1) = \frac{(a + 2\theta)}{\rho}, \quad q(\eta^*) = q(1).$$

3.3.2 Numerical example

Returning to the example from the previous section, it is once again possible to characterize the full equilibrium dynamics using numerical methods (which are documented in Appendix C). The only new parameter is α , which controls the stringency of the VaR constraint. In the baseline case I choose $\alpha = 4$, and study sensitivity of the model to changes in α in the next section. Figure 3.2 displays the results. To facilitate comparison with the previous case, I plot the new equilibrium functions alongside their unconstrained counterparts.

The economy still features the same two regimes as before, but with some important differences. First, the point of regime shift, which endogenously determines when fire sales are initiated, has shifted to the right along the state space. Second, the equilibrium price function in the constrained economy is always at or below the price function in the unconstrained economy. Third, maximal volatility in the economy is lower. These three features are common across all the various parameter configurations I consider in the next section. I will therefore spend some space discussing each point in turn.

Earlier fire sales

The endogenous point of regime switch, at which experts begin to sell off the loans they have made to finance the risky project, now occurs at a higher point on the state space. The intuition for this result is straightforward: whenever the VaR constraint binds, it means that the expert would like to engage in more intermediation, but is restricted from doing so. It should not be optimal for the expert to cut back on his intermediation activity at a point lower than in the unconstrained case, since otherwise he could have done just the same in the original economy. Of course, this type of “revealed preference” argument is only a partial equilibrium story. Since prices can change to alter economic incentives, it could be that general equilibrium effects might be strong enough to overturn this intuition. The numerical examples I consider show that this is not the case.

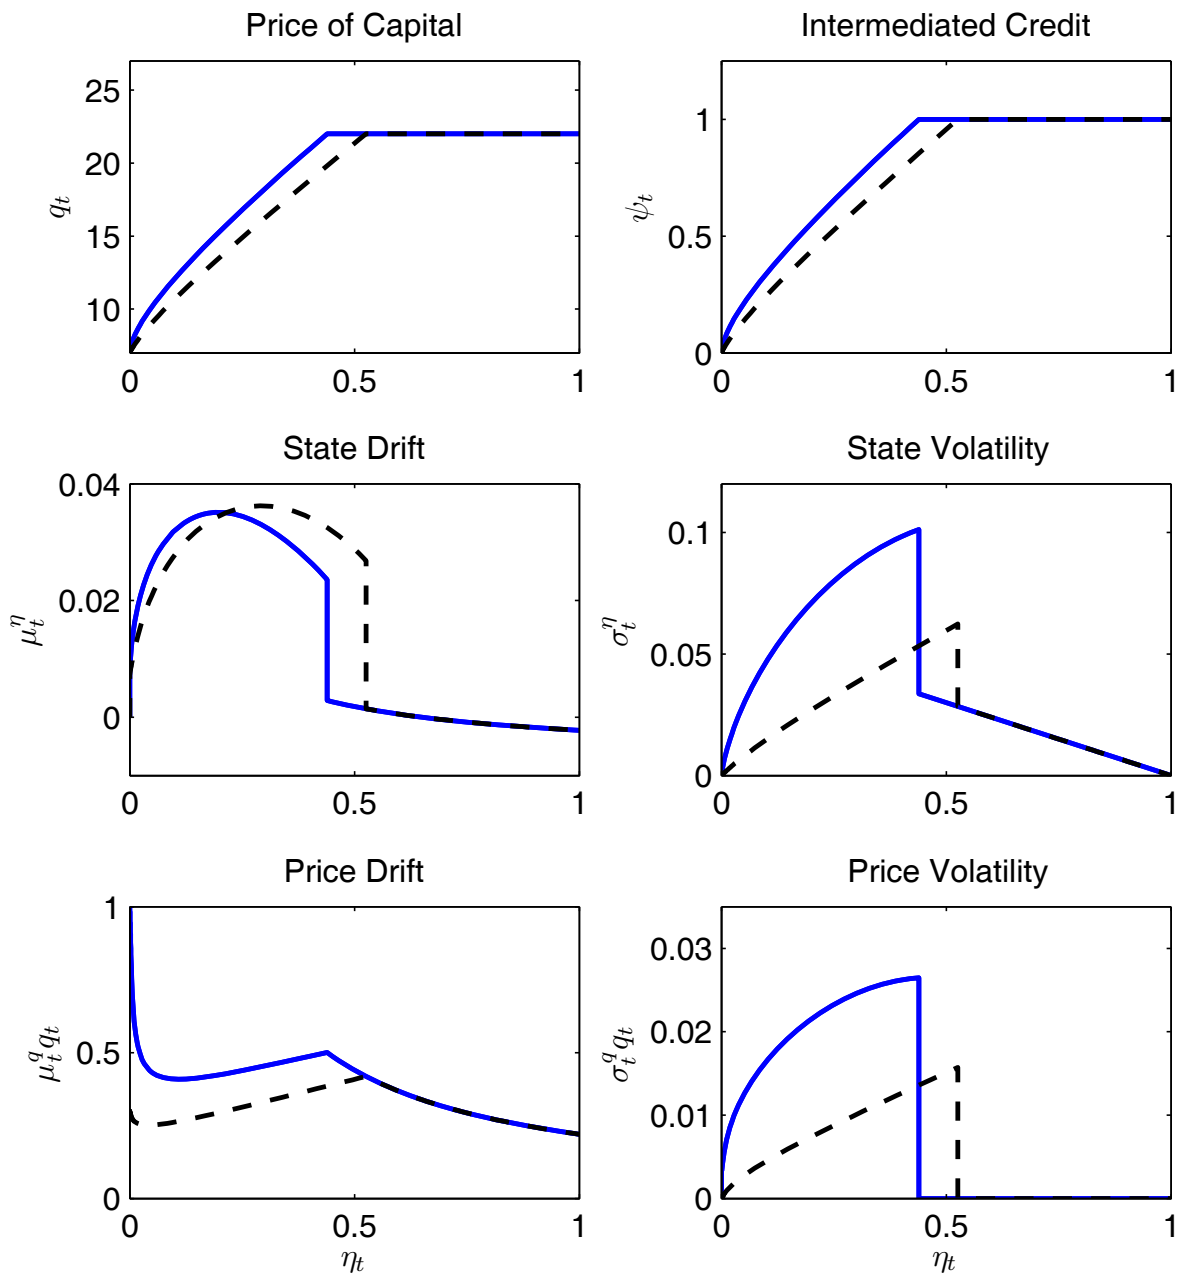


Figure 3.2: Equilibrium variables as functions of η_t in the unconstrained equilibrium (solid line) and the constrained equilibrium (dotted line).

Lower price

The price function $q_t = q(\eta_t)$ in the constrained economy is never higher than in the unconstrained one. This is closely related to the previous point. At each value on the state space, the level of expert-intermediated financing is at or below the level in the unconstrained economy. The equilibrium price reflects the marginal valuation of the risky projects weighted by the fraction of each type of agent providing financing, through the market clearing condition

$$\psi_t = \frac{\rho q_t - \underline{a} - 2\theta}{a - \underline{a}}.$$

Since the experts are constrained from extending credit, households are forced to step in, using their inefficient intermediation technology. In equilibrium, the price must fall in order to induce households to do this.

Lower volatility

Probably the most important result I find is that the VaR constraint can actually reduce volatility in the economy. However, this result depends on two offsetting effects in this economy, which can be seen from the graph. On one hand, volatility is lower along the region of the state space corresponding to the “crisis regime” in the unconstrained economy. On the other hand, volatility is higher for a small fraction of the state space where fire sales had not yet begun. To formally determine whether this second effect is enough to offset the first, it is necessary to understand what happens to the stationary distribution of the economy under a VaR constraint, which I will turn to in the following section.

This result is in line with the predictions of the similar models with preference heterogeneity considered by [Chabakauri \(2012\)](#) and [Prieto \(2013\)](#), but stands in contrast to the analysis of [Danielsson et al. \(2011\)](#). The main reason for this difference is that the third of these papers considers the special case where experts are risk neutral and maximize the instantaneous expected return on their net worth. The first two, along with the analysis

presented here, model all agents as having preferences over consumption streams, and allow for some degree of risk aversion on the behalf of experts. Formulating a dynamic continuous time model of the type considered here with risk neutral *and* risk averse agents with preferences over consumption can lead to considerable difficulties, and generally involves the tools of singular stochastic control, which are beyond the scope of this paper.

3.4 Stationary Distribution

Now that the equilibrium in each of the economies has been fully characterized, it is possible to study the stationary distribution implied in each case, and draw conclusions about what it implies for endogenous volatility. The stationary distribution of the economy represents the average location of the state variable η_t on the interval $[0, 1]$ as $t \rightarrow \infty$ when started at some initial point $\eta_0 \in (0, 1)$. Following He and Krishnamurthy (2012), I obtain this distribution in each case by numerically simulating the model. Specifically, I begin the economy at an initial state ($\eta_0 = 0.5, D_0 = 1$) and simulate the economy for 5000 years by mapping a series of independent random normal shocks dZ_t into changes in the state variable η_t using the equilibrium law of motion, $d\eta_t = \mu^\eta(\eta_t)dt + \sigma^\eta(\eta_t)dZ_t$, derived above. My simulation unit is monthly, and from these monthly observations, I consider annual averages. I then repeat this procedure 1000 times for a total of 60 million draws. I disregard the first 1000 years of data and use the remaining observations to plot the stationary distribution. I also verify that my results are not sensitive to changing the initial value η_0 , ensuring that the economy has indeed converged to its stationary limit. Further details about the simulation method are included in Appendix C. In what follows, I first describe some features of this distribution in the unconstrained case, and then compare it to the constrained one.

3.4.1 Unconstrained model

The stationary distribution is plotted in Figure 3.3. On average, the economy spends most of its time in normal states of the world, where experts are responsible for all project financing. However, the distribution is skewed to the left, which means that the economy sometimes transits to crisis states, where expert intermediation breaks down and resources are wasted. The fact that this economy has a non-degenerate stationary distribution is somewhat unique within the class of two-agent continuous-time models. A common difficulty is that in the long run, one agent (the more productive, more patient, or more risk-tolerant one) tends to dominate the aggregate resources in the economy, and the other disappears (cf. Dumas (1989) and Wang (1996)). My assumption concerning exogenous labor income helps me to avoid this problem. Intuitively, even if experts have managed to accumulate all of the wealth in the economy that comes from the risky projects, households will still not vanish, since they receive an exogenous income flow each period independent of their portfolio choices. The mathematical translation of this intuition is that the upper bound of the state space must be a reflecting barrier. To see why this is the case, observe that the second part of Proposition 2 implies that when $\psi_t = 1$,

$$d\eta_t = \left[\frac{(1 - \eta_t)^2}{\eta_t} \sigma^2 + \rho \left(\frac{a}{a + 2\theta} - 1 \right) \eta_t + \frac{\rho\theta}{a + 2\theta} \right] dt + (1 - \eta_t) \sigma dZ_t.$$

As $\eta_t \rightarrow 1$, the volatility of the state approaches zero, so that uncertainty vanishes. The drift, however, approaches a limit of

$$\rho \left(\frac{a + \theta}{a + 2\theta} - 1 \right) < 0.$$

Therefore as it approaches its upper limit, η_t becomes a deterministic process with negative drift, so it must fall back down again. On the other hand, if it were ever optimal for experts to refuse to perform any of the intermediation in the economy, so that $\psi_t = 0$, Proposition

2 implies that $q_t = \underline{q} = (\underline{a} + 2\theta)/\rho$, and therefore that the law of motion for η_t is

$$d\eta_t = \left[\frac{\eta_t^2}{(1 - \eta_t)} \sigma^2 + \rho \left(\frac{\underline{a}}{\underline{a} + 2\theta} - 1 \right) \eta_t + \frac{\rho\theta}{\underline{a} + 2\theta} \right] dt + \eta_t \sigma dZ_t.$$

As $\eta_t \rightarrow 0$, the volatility again approaches zero, but now the drift approaches a limit of

$$\frac{\rho\theta}{\underline{a} + 2\theta} > 0.$$

Therefore, η_t becomes a deterministic process with positive drift, so it must increase. Without the labor income, $\theta \rightarrow 0$, the drift would indeed converge to zero in both cases, and once η_t reached either 0 or 1, it would become absorbed there. The stationary distribution would therefore be degenerate, and introducing constraints, as I will in the next section, would have no effect on the long-run average volatility in the economy. This type of modeling device is also employed by [He and Krishnamurthy \(2013\)](#). [Brunnermeier and Sannikov \(2014\)](#) achieve stationarity instead by assuming different rates of time preference, as in [Kiyotaki and Moore \(1997\)](#), which is closely related to the constant death probability of [Bernanke et al. \(1999\)](#).

3.4.2 Constrained model

A comparison between the stationary distributions of the constrained and unconstrained economies is presented in Figure 3.2. Before discussing the results, recall from Section 3.2 that $\Theta(\eta) \equiv \{\eta \in [0, 1] \mid \eta < \eta^*\}$ denotes the set of crisis states, where η^* is the point of regime switch in the *unconstrained* economy. As is clear from the graph, the main difference between these distributions occurs in the tail; that is, in the probability that the economy finds itself on values of the state space where expert intermediation is impaired. The left tail of the constrained economy is lower. In this baseline case, I find that

$$P(\eta \in \Theta(\eta)) = 0.5869\% > 0.1517\% = \tilde{P}(\eta \in \Theta(\eta)),$$




Figure 3.3: Stationary density function $p(\eta)$, computed using simulation-based techniques.

where $P(\cdot)$ denotes probability taken with respect to the stationary distribution in the unconstrained equilibrium, and $\tilde{P}(\cdot)$ corresponds to the constrained equilibrium. The additional probability mass instead finds itself over values of η where experts are more highly capitalized; this is what accounts for the higher peak. The economy thus spends less time on average in crisis states.

This change in the distribution can be understood in the following way. First, because of the constraint, experts are forced to begin selling their assets at a higher point on the state space. This means that the adverse price spiral is initiated before the economy reaches the point η^* , that is, the crisis regime. However, remember that at higher points on the state space, experts are more highly capitalized. So, even though the same feedback effect of prices leads to further losses as before, experts now have more net worth to absorb those losses. As a result, they do not have to sell as much of the risky asset right away. Of course, this implies that there is not as severe a need for households to step in, and by market clearing the price does not need to fall by as much. The lower severity of price decline exactly translates into lower endogenous price volatility and therefore a lower volatility of net worth. This is the reason why the spike in σ_t^η at the point of regime shift shown in Figure 3.4, is much less pronounced than compared with the unconstrained economy.

The lower volatility below η^* (together with the higher drift just below η^*) translates into a lower probability of the economy staying in the crisis regime. As η falls because experts are hit by adverse fundamental shocks dZ_t , they are forced to sell off their assets earlier, but then the lower volatility and higher drift mean that the state more quickly transits out of this fire sale regime back into the normal states. Interestingly, this explanation accords nicely with the basic motivation for introducing VaR constraints: sell assets earlier, before things get really bad, so that valuable intermediaries have enough of an equity buffer to weather the storm.




Figure 3.4: Stationary density function $p(\eta)$ for the constrained and unconstrained economies. VaR constraints reduce the probability of being in a crisis regime by inducing experts to sell earlier, when they have a larger equity buffer to absorb losses.

Parameter	Baseline value	Economic meaning
ρ	0.05	rate of time preference
a	1	output rate with intermediation
\underline{a}	0.25	output rate without intermediation
g	0.05	dividend growth with intermediation
\underline{g}	0.01	dividend growth without intermediation
σ	0.06	fundamental volatility
θ	0.05	exogenous labor income
α	4	stringency of VaR constraint

Table 3.1: Baseline parameter configuration.

3.4.3 Robustness and limitations

To examine the sensitivity of my results to different assumptions on the structural parameters of the model, I repeat the entire exercise of the paper under several different specifications. The baseline parameter configuration is reported in Table 3.1, with a brief summary of each parameter’s basic economic meaning. Table 3.2 then reports the effect of introducing a VaR constraint on models with different parameter assumptions. In all cases, the volatility of the state and the probability of being in a crisis fall after introducing a VaR constraint. In some cases, I find that average price volatility can increase. This occurs whenever the stationary distribution is so compressed towards the upper boundary $\eta = 1$ that adding the constraint does not make as much of a difference, since the economy almost always features well-capitalized experts and very rarely enters crisis regimes. To the extent that crises of the sort predicted by the model are sufficiently likely events, however, I conclude that price volatility generally declines with the introduction of a VaR constraint.

With these conclusions in mind, it is also important to recognize the limitations of my analysis; I briefly highlight four here. First of all, the model does not allow for endogenous bankruptcy.¹⁴ One of the motivations for introducing VaR constraints is to protect not only against states of the world where intermediaries might become more poorly capitalized, but also against outcomes where they default on their debt obligations. It is possible that VaR

¹⁴In the model I study, no default is implicitly assumed in the agents’ optimization problems in order to ensure that they have well-defined solutions (see Dybvig and Huang (1988) on the technical importance of this assumption).

Model	$\Delta(\sigma^\eta/\eta)$	$\Delta\sigma^q$	$\Delta P(\eta \in \Theta(\eta))$
Baseline	-0.3840	-0.0306	-0.4352
$\theta = 0.005$	-0.0792	-0.0082	-0.0691
$\theta = 0.5$	-0.3608	0.0407	-0.2617
$\alpha = 2.5$	-0.1055	0.0029	-0.0851
$\alpha = 3.25$	-0.2528	-0.0209	-0.2349
$\underline{a} = 0.35$	-0.3579	-0.0126	-0.3418
$\underline{g} = 0.03$	-0.3282	-0.0265	-0.4360
$g = \underline{g} = 0.01$	-0.2175	0.0181	-0.1292
$\underline{a} = 0.35, g = \underline{g} = 0.01$	-0.1583	0.0237	-0.0847
$\theta = 0, \underline{\theta} = 0.05$	-0.7941	-0.1017	-1.1381
$\theta = 0, \underline{\theta} = 0.5$	-1.5093	-0.2379	-2.5670

Table 3.2: Sensitivity analysis. Column 2 is the change in the average percent volatility of η , column 3 is the change in the average percent volatility of risky asset prices, and column 4 is the change in probability of crisis when a VaR constraint is imposed. In all cases, the volatility of the state and the probability of being in a crisis fall after introducing a VaR constraint. In some cases, average price volatility increases.

constraints may have additional benefits associated with a reduction in bankruptcy probabilities. Incorporating this feature into the model, as in the seminal work of [Eaton and Gersovitz \(1981\)](#), would therefore supplement its testable predictions, and be particularly relevant for understanding the effects of VaR in particular. Second, I have assumed that intermediation is socially valuable in a reduced-form way. Clearly, a more careful treatment of the mechanism by which intermediation facilitates risky production decisions is also important to consider. Third, I have abstracted from many dimensions of heterogeneity it is possible to consider, mostly because they have been discussed at length in the existing literature. But expanding my analysis to incorporate those dimensions would also be a fruitful line of research. Fourth, and most importantly, I have not mentioned anything about the welfare implications of VaR. While making statements about the effects of VaR on things like average volatility is certainly interesting in its own right, one of the main themes of modern macroeconomic theory is that from a policy perspective, stabilization typically should not be understood as an end in itself. Rather, alternative policies should be evaluated based on the model-implied welfare of private agents. Unfortunately, most models which only rely on het-

erogeneous preferences or beliefs leave no room for things like VaR to be Pareto-improving. However, the model I present here does not have that difficulty, and therefore represents an important first step toward addressing this issue in a fully dynamic general equilibrium model.

3.5 Conclusion

I study the general equilibrium effects of introducing a Value-at-Risk (VaR) constraint into a dynamic continuous-time economy with homogeneous preferences, inefficient endogenous volatility, fire sales, and economically valuable financial intermediation. The model is populated by two types of agents: households, who can only inefficiently finance risky projects, and experts, who have access to an efficient intermediation technology. I find that through its impact on the stationary distribution of wealth in the economy, a VaR constraint can reduce the average level of endogenous volatility, and lower the probability of entering a crisis regime. It does so by forcing agents to sell off their asset holdings earlier than they otherwise would, while they have a larger equity buffer to absorb losses. This chapter is the first study to explore the effects of a VaR constraint in a model that does not feature any heterogeneity in preferences or beliefs, and in which endogenous volatility and crises are socially inefficient.

Bibliography

- Abraham, Katharine G. and Lawrence F. Katz (1986) “Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?,” *Journal of Political Economy*, 94(3):507–522.
- Adam, Klaus and Michael Woodford (2012) “Robustly optimal monetary policy in a micro-founded New Keynesian model,” *Journal of Monetary Economics*, 59(5):468–487.
- Akerlof, George A. and Robert J. Shiller (2009) *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism*, Princeton: Princeton University Press.
- Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Lindé (2011) “Firm-specific capital, nominal rigidities and the business cycle,” *Review of Economic Dynamics*, 14(2):225–247.
- An, Sungbae and Frank Schorfheide (2007) “Bayesian Analysis of DSGE Models,” *Econometric Reviews*, 26(2-4):113–172.
- Andersen, Torben G., Tim Bollerslev, Peter F. Christoffersen, and Francis X. Diebold (2013) “Chapter 17 - Financial Risk Measurement for Financial Risk Management,” in M. H. George M. Constantinides and R. M. Stulz (Eds.), *Handbook of the Economics of Finance*, Volume 2, Part B, pp. 1127 – 1220, Chicago: Elsevier.
- Angeletos, George-Marios, Fabrice Collard, and Harris Dellas (2014) “Quantifying Confidence,” Working Paper 20807, National Bureau of Economic Research.

- Angeletos, George-Marios and Jennifer La'O (2010) "Noisy Business Cycles," in *NBER Macroeconomics Annual 2009*, Volume 24, pp. 319–378, Chicago: University of Chicago Press.
- Arellano, Cristina, Yan Bai, and Patrick J. Kehoe (2012) "Financial frictions and fluctuations in volatility," Staff Report 466, Federal Reserve Bank of Minneapolis.
- Bachmann, Rüdiger and Christian Bayer (2011) "Uncertainty Business Cycles - Really?," Working Paper 16862, National Bureau of Economic Research.
- Bachmann, Rüdiger, Steffen Elstner, and Eric R. Sims (2013) "Uncertainty and Economic Activity: Evidence from Business Survey Data," *American Economic Journal: Macroeconomics*, 5(2):217–49.
- Bai, Jushan and Serena Ng (2002) "Determining the Number of Factors in Approximate Factor Models," *Econometrica*, 70(1):191–221.
- Bai, Jushan and Serena Ng (2004) "A PANIC Attack on Unit Roots and Cointegration," *Econometrica*, 72(4):1127–1177.
- Bai, Jushan and Serena Ng (2006) "Confidence Intervals for Diffusion Index Forecasts and Inference for Factor-Augmented Regressions," *Econometrica*, 74(4):1133–1150.
- Bai, Jushan and Serena Ng (2008) "Forecasting economic time series using targeted predictors," *Journal of Econometrics*, 146(2):304–317.
- Baker, Scott R. and Nicholas Bloom (2013) "Does Uncertainty Reduce Growth? Using Disasters as Natural Experiments," Working Paper 19475, National Bureau of Economic Research.
- Baker, Scott R., Nicholas Bloom, and Steven J. Davis (2013) "Measuring Economic Policy Uncertainty," Working Paper, Chicago Booth.

- Bansal, Ravi, Dana Kiku, and Amir Yaron (2007) “Risks for the Long Run: Estimation and Inference,” Working Paper, Duke University.
- Bansal, Ravi, Dana Kiku, and Amir Yaron (2012) “An Empirical Evaluation of the Long-Run Risks Model for Asset Prices,” *Critical Finance Review*, 1(1):183–221.
- Bansal, Ravi and Amir Yaron (2004) “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *The Journal of Finance*, 59(4):1481–1509.
- Barsky, Robert B., Susanto Basu, and Keyoung Lee (2014) “Whither News Shocks?,” in *NBER Macroeconomics Annual 2014*, Volume 29, Chicago: University of Chicago Press.
- Barsky, Robert B. and Eric R. Sims (2011) “News shocks and business cycles,” *Journal of Monetary Economics*, 58(3):273–289.
- Barsky, Robert B. and Eric R. Sims (2012) “Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence,” *American Economic Review*, 102(4):1343–77.
- Basak, Suleyman and Domenico Cuoco (1998) “An equilibrium model with restricted stock market participation,” *Review of Financial Studies*, 11(2):309–341.
- Basak, S and A Shapiro (2001) “Value-at-risk-based risk management: optimal policies and asset prices,” *Review of Financial Studies*, 14(2):371–405.
- Basu, Susanto and Brent Bundick (2012) “Uncertainty Shocks in a Model of Effective Demand,” Working Paper 18420, National Bureau of Economic Research.
- Beaudry, Paul and Bernd Lucke (2010) “Letting Different Views about Business Cycles Compete,” in *NBER Macroeconomics Annual 2009*, Volume 24, pp. 413–455, Chicago: University of Chicago Press.
- Beaudry, Paul and Franck Portier (2006) “Stock Prices, News, and Economic Fluctuations,” *American Economic Review*, 96(4):1293–1307.

- Beaudry, Paul and Franck Portier (2014) “News-Driven Business Cycles: Insights and Challenges,” *Journal of Economic Literature*, 52(4):993–1074.
- Bekaert, Geert, Marie Hoerova, and Marco Lo Duca (2013) “Risk, uncertainty and monetary policy,” *Journal of Monetary Economics*, 60(7):771–788.
- Bernanke, Ben S., Mark Gertler, and Simon Gilchrist (1999) “The financial accelerator in a quantitative business cycle framework,” in J. B. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, Volume 1, Part C, Chapter 21, pp. 1341–1393, Amsterdam: Elsevier.
- Bidder, R.M. and M.E. Smith (2012) “Robust animal spirits,” *Journal of Monetary Economics*, 59(8):738–750.
- Blanchard, Olivier J., Jean-Paul L’Huillier, and Guido Lorenzoni (2013) “News, Noise, and Fluctuations: An Empirical Exploration,” *American Economic Review*, 103(7):3045–70.
- Bloom, Nicholas (2009) “The Impact of Uncertainty Shocks,” *Econometrica*, 77(3):623–685.
- Bloom, Nick, Stephen Bond, and John Van Reenen (2007) “Uncertainty and Investment Dynamics,” *The Review of Economic Studies*, 74(2):391–415.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry (2012) “Really Uncertain Business Cycles,” Working Paper 18245, National Bureau of Economic Research.
- Boldrin, Michele, Lawrence J. Christiano, and Jonas D. M. Fisher (2001) “Habit Persistence, Asset Returns, and the Business Cycle,” *American Economic Review*, 91(1):149–166.
- Borovička, Jaroslav, Lars P. Hansen, and José A. Scheinkman (2014) “Misspecified Recovery,” Working Paper 20209, National Bureau of Economic Research.
- Brunnermeier, Markus K. and Lasse Heje Pedersen (2009) “Market Liquidity and Funding Liquidity,” *Review of Financial Studies*, 22(6):2201–2238.

- Brunnermeier, Markus K. and Yuliy Sannikov (2014) “A Macroeconomic Model with a Financial Sector,” *American Economic Review*, 104(2):379–421.
- Cagetti, Marco, Lars Peter Hansen, Thomas Sargent, and Noah Williams (2002) “Robustness and Pricing with Uncertain Growth,” *Review of Financial Studies*, 15(2):363–404.
- Campbell, John Y. and John H. Cochrane (1999) “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107(2):205–251.
- Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2012) “Common drifting volatility in large Bayesian VARs,” Working Paper 1206, Federal Reserve Bank of Cleveland.
- Carrillo, Julio, Patrick Fève, and Julien Matheron (2007) “Monetary Policy Inertia or Persistent Shocks: A DSGE Analysis,” *International Journal of Central Banking*, 3(2):1–38.
- Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark (2000) “Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to Be True?,” *American Economic Review*, 90(4):787–805.
- Chabakauri, Georgy (2012) “Asset Pricing with Heterogeneous Investors and Portfolio Constraints,” Meeting Paper 636, Society for Economic Dynamics.
- Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan (2009) “New Keynesian Models: Not Yet Useful for Policy Analysis,” *American Economic Journal: Macroeconomics*, 1(1):242–66.
- Christiano, Lawrence, Cosmin Ilut, Roberto Motto, and Massimo Rostagno (2010) “Monetary policy and stock market booms,” in *Macroeconomic Challenges: the Decade Ahead*, Economic Policy Symposium, Jackson Hole, pp. 85–145, Federal Reserve Bank of Kansas City.

- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005) “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113(1):1–45.
- Cochrane, John H. (1994) “Permanent and Transitory Components of GNP and Stock Prices,” *The Quarterly Journal of Economics*, 109(1):241–265.
- Cochrane, John H. (2005) *Asset Pricing*, Princeton: Princeton University Press.
- Cogley, Timothy and Thomas J. Sargent (2005) “Drifts and volatilities: monetary policies and outcomes in the post {WWII} {US},” *Review of Economic Dynamics*, 8(2):262–302
Monetary Policy and Learning.
- Croce, Mariano Massimiliano (2014) “Long-run productivity risk: A new hope for production-based asset pricing?,” *Journal of Monetary Economics*, 66(0):13–31.
- Cuoco, Domenico and Hua He (2001) “Dynamic Aggregation and Computation of Equilibria in Finite-Dimensional Economies with Incomplete Financial Markets,” *Annals of Economics and Finance*, 2(2):265–296.
- Cuoco, Domenico, Hua He, and Sergei Isaenko (2008) “Optimal Dynamic Trading Strategies with Risk Limits,” *Operations Research*, 56(2):358–368.
- Cúrdia, Vasco and Ricardo Reis (2010) “Correlated Disturbances and U.S. Business Cycles,” Working Paper 15774, National Bureau of Economic Research.
- Dai, Qiang and Kenneth J. Singleton (2000) “Specification Analysis of Affine Term Structure Models,” *The Journal of Finance*, 55(5):1943–1978.
- D’Amico, Stefania and Athanasios Orphanides (2008) “Uncertainty and disagreement in economic forecasting,” Finance and Economics Discussion Series 2008-56, Board of Governors of the Federal Reserve System (U.S.).

- Daniélsson, Jón, Paul Embrechts, Charles Goodhart, Con Keating, Felix Muennich, Olivier Renault, and Hyun Song Shin (2001) “An Academic Response to Basel II,” Special Paper 130, London School of Economics Financial Markets Group.
- Daniélsson, Jón, Hyun Song Shin, and Jean-Pierre Zigrand (2004) “The impact of risk regulation on price dynamics,” *Journal of Banking & Finance*, 28(5):1069–1087.
- Daniélsson, Jón, Hyun Song Shin, and Jean-Pierre Zigrand (2011) “Balance Sheet Capacity and Endogenous Risk,” FMG Discussion Papers dp665, Financial Markets Group.
- Del Negro, Marco and Stefano Eusepi (2011) “Fitting observed inflation expectations,” *Journal of Economic Dynamics and Control*, 35(12):2105 – 2131.
- Diether, Karl B., Christopher J. Malloy, and Anna Scherbina (2002) “Differences of Opinion and the Cross Section of Stock Returns,” *The Journal of Finance*, 57(5):2113–2141.
- Duffie, Darrell and Jun Pan (1997) “An Overview of Value at Risk,” *The Journal of Derivatives*, 4(3):7–49.
- Dumas, B (1989) “Two-person dynamic equilibrium in the capital market,” *Review of Financial Studies*, 2(2):157–188.
- Durbin, James and Siem Jan Koopman (2012) *Time Series Analysis by State Space Methods* (2 ed.), Oxford: Oxford University Press.
- Dybvig, Philip H. and Chi-fu Huang (1988) “Nonnegative Wealth, Absence of Arbitrage, and Feasible Consumption Plans,” *The Review of Financial Studies*, 1(4):377–401.
- Eaton, Jonathan and Mark Gersovitz (1981) “Debt with Potential Repudiation: Theoretical and Empirical Analysis,” *The Review of Economic Studies*, 48(2):289–309.
- Epstein, Larry G. and Stanley E. Zin (1989) “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57(4):937–969.

- Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin (2000) “Optimal monetary policy with staggered wage and price contracts,” *Journal of Monetary Economics*, 46(2):281–313.
- Eusepi, Stefano and Bruce Preston (2011) “Expectations, Learning, and Business Cycle Fluctuations,” *American Economic Review*, 101(6):2844–72.
- Fama, Eugene F. and Kenneth R. French (1992) “The Cross-Section of Expected Stock Returns,” *The Journal of Finance*, 47(2):427–465.
- Fama, Eugene F. and Kenneth R. French (1993) “Common risk factors in the returns on stocks and bonds,” *Journal of Financial Economics*, 33(1):3–56.
- Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Juan F. Rubio-Ramírez, and Martín Uribe (2011) “Risk Matters: The Real Effects of Volatility Shocks,” *American Economic Review*, 101(6):2530–61.
- Gârleanu, Nicolae and Stavros Panageas (2015) “Young, Old, Conservative and Bold. The implication of finite lives and heterogeneity for Asset Pricing,” *Journal of Political Economy*, (Forthcoming).
- Gârleanu, Nicolae and Lasse Heje Pedersen (2011) “Margin-based Asset Pricing and Deviations from the Law of One Price,” *Review of Financial Studies*, 24(6):1980–2022.
- Gilchrist, Simon, Jae W. Sim, and Egon Zakrajsek (2014) “Uncertainty, Financial Frictions, and Investment Dynamics,” Finance and Economics Discussion Series 2014-69, Board of Governors of the Federal Reserve System.
- Gomme, Paul and Paul Klein (2011) “Second-order approximation of dynamic models without the use of tensors,” *Journal of Economic Dynamics and Control*, 35(4):604–615.
- Gourio, François (2012) “Disaster Risk and Business Cycles,” *American Economic Review*, 102(6):2734–66.

- Greasly, David and Jakob B. Madsen (2006) "Investment and Uncertainty: Precipitating the Great Depression in the United States," *Economica*, 73(291):393–412.
- Greenspan, Alan (2003) "Testimony Before the Committee on Financial Services, U.S. House of Representatives," Federal Reserve Board's semiannual monetary policy report to the Congress, July 15, 2003.
- Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman (1988) "Investment, Capacity Utilization, and the Real Business Cycle," *The American Economic Review*, 78(3):pp. 402–417.
- Hall, Robert E. (1978) "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86(6):pp. 971–987.
- Hansen, Lars Peter and Thomas J. Sargent (2005) "Robust estimation and control under commitment," *Journal of Economic Theory*, 124(2):258–301.
- Hansen, Lars Peter and Thomas J. Sargent (2007) "Recursive robust estimation and control without commitment," *Journal of Economic Theory*, 136(1):1–27.
- Hansen, Lars Peter, Thomas J. Sargent, and Thomas D. Tallarini (1999) "Robust Permanent Income and Pricing," *The Review of Economic Studies*, 66(4):873–907.
- Hassan, Tarek A. and Thomas M. Mertens (2014) "Information Aggregation in a DSGE Model," in *NBER Macroeconomics Annual 2014*, Volume 29, Chicago: University of Chicago Press (Forthcoming).
- Hassler, John (2001) "Uncertainty and the Timing of Automobile Purchases," *Scandinavian Journal of Economics*, 103(2):351–66.
- He, Zhiguo and Arvind Krishnamurthy (2012) "A Model of Capital and Crises," *The Review of Economic Studies*, 79(2):735–777.

- He, Zhiguo and Arvind Krishnamurthy (2013) “Intermediary Asset Pricing,” *American Economic Review*, 103(2):732–70.
- Herskovic, Bernard, Bryan T. Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh (2014) “The Common Factor in Idiosyncratic Volatility: Quantitative Asset Pricing Implications,” Working Paper 20076, National Bureau of Economic Research.
- Hodrick, Robert J and Edward C Prescott (1997) “Postwar U.S. Business Cycles: An Empirical Investigation,” *Journal of Money, Credit and Banking*, 29(1):1–16.
- Ilut, Cosmin L. and Martin Schneider (2014) “Ambiguous Business Cycles,” *American Economic Review*, 104(8):2368–99.
- Jaimovich, Nir and Sergio Rebelo (2009) “Can News about the Future Drive the Business Cycle?,” *American Economic Review*, 99(4):1097–1118.
- Jeanne, Olivier and Anton Korinek (2010) “Managing Credit Booms and Busts: A Pigouvian Taxation Approach,” Working Paper 16377, National Bureau of Economic Research.
- Jermann, Urban J. (1998) “Asset pricing in production economies,” *Journal of Monetary Economics*, 41(2):257–275.
- Jinnai, Ryo (2013) “News shocks and inflation,” *Economics Letters*, 119(2):176–179.
- Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti (2011) “Investment shocks and the relative price of investment,” *Review of Economic Dynamics*, 14(1):102–121.
- Kaltenbrunner, Georg and Lars A. Lochstoer (2010) “Long-Run Risk through Consumption Smoothing,” *Review of Financial Studies*, 23(8):3190–3224.
- Kastner, Gregor and Sylvia Frühwirth-Schnatter (2014) “Ancillarity-sufficiency interweaving strategy (ASIS) for boosting {MCMC} estimation of stochastic volatility models,”

Computational Statistics & Data Analysis, 76(0):408–423 CFEnetwork: The Annals of Computational and Financial Econometrics 2nd Issue.

Kim, Sangjoon, Neil Shephard, and Siddhartha Chib (1998) “Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models,” *The Review of Economic Studies*, 65(3):361–393.

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo (1988) “Production, growth and business cycles: I. The basic neoclassical model,” *Journal of Monetary Economics*, 21(23):195–232.

Kiyotaki, Nobuhiro and John Moore (1997) “Credit Cycles,” *Journal of Political Economy*, 105(2):211–248.

Klein, Paul (2000) “Using the generalized Schur form to solve a multivariate linear rational expectations model,” *Journal of Economic Dynamics and Control*, 24(10):1405–1423.

Knotek, Edward S. II. and Shujaat Khan (2011) “How do households respond to uncertainty shocks?,” *Federal Reserve Bank of Kansas City Economic Review*, 96(2):5–34.

Kreps, David M. and Evan L. Porteus (1978) “Temporal Resolution of Uncertainty and Dynamic Choice Theory,” *Econometrica*, 46(1):185–200.

Krusell, Per and Anthony A. Smith, Jr. (1998) “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5):867–896.

Lahiri, Kajal and Xuguang Sheng (2010) “Measuring forecast uncertainty by disagreement: The missing link,” *Journal of Applied Econometrics*, 25(4):514–538.

Le, Anh, Kenneth J. Singleton, and Qiang Dai (2010) “Discrete-Time Affine^Q Term Structure Models with Generalized Market Prices of Risk,” *Review of Financial Studies*, 23(5):2184–2227.

- Leahy, John V. and Toni M. Whited (1996) “The Effect of Uncertainty on Investment: Some Stylized Facts,” *Journal of Money, Credit and Banking*, 28(1):64–83.
- Leduc, Sylvain and Zheng Liu (2014) “Uncertainty shocks are aggregate demand shocks,” Working Paper Series 2012-10, Federal Reserve Bank of San Francisco.
- Lettau, Martin and Sydney C. Ludvigson (2004) “Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption,” *American Economic Review*, 94(1):276–299.
- Li, Erica X.N. and Francisco Palomino (2014) “Nominal rigidities, asset returns, and monetary policy,” *Journal of Monetary Economics*, 66(0):210–225.
- Longstaff, Francis A. and Jiang Wang (2012) “Asset Pricing and the Credit Market,” *Review of Financial Studies*, 25(11):3169–3215.
- Lorenzoni, Guido (2008) “Inefficient Credit Booms,” *The Review of Economic Studies*, 75(3):809–833.
- Lorenzoni, Guido (2009) “A Theory of Demand Shocks,” *American Economic Review*, 99(5):2050–84.
- Ludvigson, Sydney C. and Serena Ng (2007) “The empirical riskreturn relation: A factor analysis approach,” *Journal of Financial Economics*, 83(1):171–222.
- Ludvigson, Sydney C. and Serena Ng (2009) “Macro Factors in Bond Risk Premia,” *Review of Financial Studies*, 22(12):5027–5067.
- Ludvigson, Sydney C. and Serena Ng (2010) “A Factor Analysis of Bond Risk Premia,” in A. Ullah and D. E. A. Giles (Eds.), *Handbook of Empirical Economics and Finance*, Volume 1, Chapter 11, pp. 313–372, Boca Raton: Chapman and Hall.
- Malkhozov, Aytek (2014) “Asset prices in affine real business cycle models,” *Journal of Economic Dynamics and Control*, 45(0):180–193.

- Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers (2003) “Disagreement about Inflation Expectations,” in *NBER Macroeconomics Annual 2003*, Volume 18, pp. 209–248, Chicago: The University of Chicago Press.
- Mehra, Rajnish and Edward C. Prescott (1985) “The equity premium: A puzzle,” *Journal of Monetary Economics*, 15(2):145–161.
- Mele, Antonio (2007) “Asymmetric stock market volatility and the cyclical behavior of expected returns,” *Journal of Financial Economics*, 86(2):446–478.
- Merton, Robert C. (1969) “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case,” *The Review of Economics and Statistics*, 51(3):247–257.
- Merton, Robert C. (1971) “Optimum consumption and portfolio rules in a continuous-time model,” *Journal of Economic Theory*, 3(4):373–413.
- Milani, Fabio (2011) “Expectation Shocks and Learning as Drivers of the Business Cycle*,” *The Economic Journal*, 121(552):379–401.
- Miyamoto, Wataru and Thuy Lan Nguyen (2014) “News Shocks and Business Cycles: Evidence from Forecast Data,” Working Paper, Columbia University.
- Nakamura, Emi, Dmitriy Sergeyev, and Jón Steinsson (2012) “Growth-Rate and Uncertainty Shocks in Consumption: Cross-Country Evidence,” Working Paper 18128, National Bureau of Economic Research.
- Orlik, Anna and Laura Veldkamp (2014) “Understanding Uncertainty Shocks and the Role of Black Swans,” Working Paper 20445, National Bureau of Economic Research.
- Pagan, Adrian (1984) “Econometric Issues in the Analysis of Regressions with Generated Regressors,” *International Economic Review*, 25(1):221–247.
- Piazzesi, Monika, Juliana Salomao, and Martin Schneider (2013) “Trend and Cycle in Bond Premia,” Working Paper, Stanford University.

- Pigou, Arthur C. (1927) *Industrial Fluctuations*, London: Macmillan.
- Preston, Bruce (2005) “Learning about Monetary Policy Rules when Long-Horizon Expectations Matter,” *International Journal of Central Banking*, 1(2):81–126.
- Prieto, Rodolfo (2013) “Dynamic equilibrium with heterogeneous agents and risk constraints,” Working Paper, Boston University.
- Rauh, Joshua D. and Amir Sufi (2012) “Explaining Corporate Capital Structure: Product Markets, Leases, and Asset Similarity,” *Review of Finance*, 16(1):115–155.
- Restoy, Fernando and G. Michael Rockinger (1994) “On Stock Market Returns and Returns on Investment,” *The Journal of Finance*, 49(2):543–556.
- Romer, Christina D. (1990) “The Great Crash and the Onset of the Great Depression,” *The Quarterly Journal of Economics*, 105(3):597–624.
- Rudebusch, Glenn D. (2002) “Term structure evidence on interest rate smoothing and monetary policy inertia,” *Journal of Monetary Economics*, 49(6):1161–1187.
- Rudebusch, Glenn D. (2006) “Monetary Policy Inertia: Fact or Fiction?,” *International Journal of Central Banking*, 2(4):85–135.
- Rudebusch, Glenn D. and Eric T. Swanson (2012) “The Bond Premium in a DSGE Model with Long-Run Real and Nominal Risks,” *American Economic Journal: Macroeconomics*, 4(1):105–43.
- Sannikov, Yuliy (2013) “Solving Heterogenous-Agent Models with Financial Frictions: A Continuous Time Approach,” Lecture Notes, Princeton Initiative: Money, Macro, and Finance, Princeton University.
- Schaal, Edouard (2011) “Uncertainty, Productivity and Unemployment in the Great Depression,” Meeting Paper 1450, Society for Economic Dynamics.

- Schmitt-Grohé, Stephanie and Martín Uribe (2004) “Solving dynamic general equilibrium models using a second-order approximation to the policy function,” *Journal of Economic Dynamics and Control*, 28(4):755–775.
- Schmitt-Grohé, Stephanie and Martín Uribe (2006) “Comparing Two Variants of Calvo-Type Wage Stickiness,” Working Paper 12740, National Bureau of Economic Research.
- Schmitt-Grohé, Stephanie and Martín Uribe (2012) “What’s News in Business Cycles,” *Econometrica*, 80(6):2733–2764.
- Scotti, Chiara (2013) “Surprise and uncertainty indexes: real-time aggregation of real-activity macro surprises,” International Finance Discussion Paper 1093, Board of Governors of the Federal Reserve System.
- Shiller, Robert J. (2015) *Irrational Exuberance* (3 ed.), Princeton: Princeton University Press.
- Shimer, Robert (2009) “Convergence in Macroeconomics: The Labor Wedge,” *American Economic Journal: Macroeconomics*, 1(1):280–97.
- Shin, Hyun Song (2010) *Risk and Liquidity*, Oxford: Oxford University Press.
- Smets, Frank and Rafael Wouters (2007) “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97(3):586–606.
- So, Eric C. (2013) “A new approach to predicting analyst forecast errors: Do investors overweight analyst forecasts?,” *Journal of Financial Economics*, 108(3):615–640.
- Stock, James H. and Mark W. Watson (2002a) “Forecasting Using Principal Components from a Large Number of Predictors,” *Journal of the American Statistical Association*, 97(460):1167–1179.
- Stock, James H and Mark W Watson (2002b) “Macroeconomic Forecasting Using Diffusion Indexes,” *Journal of Business & Economic Statistics*, 20(2):147–162.

- Stock, James H. and Mark W. Watson (2006) “Forecasting with Many Predictors,” in C. G. G. Elliott and A. Timmermann (Eds.), *Handbook of Economic Forecasting*, Volume 1, Chapter 10, pp. 515–554, Amsterdam: Elsevier.
- Wang, Jiang (1996) “The term structure of interest rates in a pure exchange economy with heterogeneous investors,” *Journal of Financial Economics*, 41(1):75–110.
- Weil, Philippe (1989) “The equity premium puzzle and the risk-free rate puzzle,” *Journal of Monetary Economics*, 24(3):401–421.
- Woodford, Michael (2010) “Robustly Optimal Monetary Policy with Near-Rational Expectations,” *The American Economic Review*, 100(1):274–303.

Appendix A

Appendix for Chapter 1

A.1 Data Description

All data series used in estimation and their sources are listed here. All series from the BEA were downloaded from www.bea.gov, and those from the BLS from www.bls.gov, and all those from the Federal Reserve Bank of Philadelphia from www.philadelphiafed.org.

A.1.1 Raw Data

- (1) Gross Domestic Product: NIPA table 1.1.5, line 1.
- (2) Personal Consumption Expenditure on Durable Goods: NIPA table 1.1.5, line 4.
- (3) Personal Consumption Expenditure on Nondurable Goods: NIPA table 1.1.5, line 5.
- (4) Personal Consumption Expenditure on Services: NIPA table 1.1.5, line 6.
- (5) Nonresidential Gross Private Domestic Fixed Investment: NIPA table 1.1.5, line 9.
- (6) Residential Gross Private Domestic Fixed Investment: NIPA table 1.1.5, line 13.
- (7) Government Consumption Expenditure: NIPA table 3.9.5, line 2.
- (8) Government Gross Investment: NIPA table 3.9.5, line 3.

- (9) Implicit Deflator for Gross Domestic Product: NIPA table 1.1.9, line 1.
- (10) Implicit Deflator for Personal Consumption Expenditure on Durable Goods: NIPA table 1.1.9, line 4.
- (11) Implicit Deflator for Nonresidential Gross Private Domestic Fixed Investment: NIPA table 1.1.9, line 9.
- (12) Implicit Deflator for Residential Gross Private Domestic Fixed Investment: NIPA table 1.1.9, line 13.
- (13) Civilian Noninstitutional Population Over 16: BLS series LNU000000000Q.
- (14) Nonfarm Business Hours Worked: BLS series PRS85006033.
- (15) Nonfarm Business Hourly Compensation: BLS series PRS85006103.
- (14) Nonfarm Business Utilization-Adjusted Total Factor Productivity: from John Fernald. Used in Basu, Fernald, and Kimball (2006). Downloaded from www.frbsf.org.
- (15) 3 Month Treasury Bill Secondary Market Rate: FRB H.15 release. Downloaded from www.federalreserve.gov.
- (16) Stock Return Index Including Distributions: CRSP NYSE/NASDAQ/AMEX monthly value-weighted stock return index, series name VWRETD. Downloaded from CRSP through WRDS at wrds-web.wharton.upenn.edu/wrds.
- (17) Mean Forecasts of Real Gross Domestic Product for the current quarter: Federal Reserve Bank of Philadelphia Survey of Professional Forecasters.
- (18) Mean Forecasts of Real Gross Domestic Product one quarter ahead: Federal Reserve Bank of Philadelphia Survey of Professional Forecasters.
- (19) Mean Forecasts of Implicit Deflator for Gross Domestic Product for the current quarter: Federal Reserve Bank of Philadelphia Survey of Professional Forecasters.

- (20) Mean Forecasts of Implicit Deflator for Gross Domestic Product one quarter ahead: Federal Reserve Bank of Philadelphia Survey of Professional Forecasters.
- (21) Mean Forecasts of 3-Month Treasury Bill Rate one quarter ahead: Federal Reserve Bank of Philadelphia Survey of Professional Forecasters.

A.1.2 Data Transformations

- (22) Real per capita GDP: $Y_t^d = (1)/(9)/(13)$.
- (23) Real per capita consumption: $C_t = [(3) + (4)]/(9)/(13)$.
- (24) Real per capita investment: $X_t I_t = [(2) + (5) + (6)]/(9)/(13)$.
- (25) Per capita hours worked: $H_t = (14)/(13)$.
- (26) Real wage: $W_t = (15)/(9)$.
- (27) Inflation: $\Pi_t = (9)_t/(9)_{t-1}$.
- (28) Real stock market return: $R_t^e = (1 + (16))/(27)$.
- (29) Nominal risk-free return: $R_t = [1 - \frac{91}{36000}(15)]^{-1}$.
- (30) Expected per capita real GDP growth: $E_t^{\mathbb{Q}}[\Delta \ln(Y_{t+1}^d)] = \ln((18)/(17)/(13))$.
- (31) Expected inflation: $E_t^{\mathbb{Q}}[\Delta \Pi_{t+1}] = \ln((20)/(19))$.
- (32) Expected nominal risk-free return: $E_t^{\mathbb{Q}}[\ln(R_{t+1})] = -\ln([1 - \frac{91}{36000}(21)])$.
- (33) Real per capita government consumption: $G_t = [(7) + (8)]/(9)/(13)$.
- (34) Total factor productivity: $TFP_t = 1 + (14)/100$.
- (35) Real price of investment X_t = Fisher index of (10),(11),(12) using (2),(5), and (6).

A.2 Related Types of Disturbances

This section presents three different sources of non-fundamental variation that have recently been explored in the literature and clarify how those are related to the type of belief distortions modeled in this chapter. These approaches are all similar and highly complementary; nevertheless, they are theoretically distinct. I show that to empirically differentiate them, it is important to incorporate observations of agents' subjective forecasts.

A.2.1 Ambiguity

Recent work by [Ilut and Schneider \(2014\)](#) has emphasized the importance of time-varying ambiguity together with ambiguity-averse agents for explaining business cycle fluctuations. To describe a situation similar to the one in their chapter, consider a simple economic model that relates an agent's action x to his expectation of a fundamental random variable z :

$$x = \tilde{E}[z].$$

Here \tilde{E} represents the agent's subjective expectation; I use a tilde rather than a \mathbb{Q} to allow for the possibility that this expectation may not correspond to any underlying well-defined probability measure.

At the time of his action, the agent believes that z will be drawn from an interval $[b - \xi, b + \xi]$ with $\xi \geq 0$. For simplicity, he believes z will be drawn according to a symmetric probability distribution. However, he is averse to ambiguity, and bases his decisions only on the worst possible outcome in this set. So he chooses

$$x = b - \xi.$$

If instead the agent were to act based on his average belief, he would choose

$$x = b.$$

With only observations of x , it is not possible to tell these alternative hypotheses apart. However, with additional observations of the agent's average or worst-case forecast (or both), it is possible. This is because the actions of an ambiguity-averse agent should be more highly correlated with his worst-case forecast than his average forecast. The focus of this chapter is on belief distortions of the second type. This is consistent with the way that the average survey forecasts are interpreted and used in the empirical analysis.

A.2.2 Confidence

Research in progress by [Angeletos et al. \(2014\)](#) introduces exogenous variation in agents' higher-order beliefs. In settings with strategic interaction, these types of disturbances affect endogenous outcomes. For example, consider a simple economic model that relates an agent's action x_i to his expectation of a fundamental random variable z and his expectation of the aggregate action taken in the economy, $x \equiv \int x_i di$:

$$x_i = \alpha \tilde{E}_i[z] + (1 - \alpha) \tilde{E}_i[x].$$

Suppose that each atomistic agent believes that z has a mean of b . But at the same time, he maintains that all other agents in the economy believe that z has mean $b + \xi$. All agents agree to disagree about these beliefs. The random variable ξ is what the authors interpret as “confidence.” By integrating the equilibrium condition above and iterating across all higher-order beliefs, it is possible to show that the optimal aggregate action satisfies

$$x = b + \frac{(1 - \alpha)^2}{\alpha} \xi.$$

With only observations of x it is not possible to distinguish b from ξ . However, with additional observations of the average (first-order) forecast, it is possible. This is because distortions in higher-order beliefs do not affect first-order beliefs. The focus of this chapter is on belief distortions of the second type. Again, this is consistent with the way that average survey forecasts are interpreted and used in the empirical analysis.

A.2.3 Noise Traders

[Hassan and Mertens \(2014\)](#) explore the role of exogenous non-fundamental variation driven by “noise traders:” agents who randomly demand an exogenous quantity of real resources each period. As a simple example of how noise traders can look like belief distortions, consider an economy with two agents. The first has mean-variance preferences with a unit coefficient of absolute risk aversion and maximizes his expected utility of terminal wealth. The other is a noise trader who exogenously demands an amount ξ of the resource. The resource is traded competitively at price x and pays of a random amount z in the next period. The demand function of the utility-maximizing agent depends on his expectation of the terminal payout of that resource and its price. By market clearing:

$$1 = \frac{\tilde{E}[z] - x}{\tilde{\text{var}}[z]} + \xi.$$

Finally, suppose that the utility-maximizing agent believes that z has a mean of $(1 + b)$ and a variance of one. Then solving for the equilibrium price,

$$x = b + \xi.$$

With only observations of x it is not possible to distinguish b from ξ . However, with additional observations of the average forecast $\tilde{E}[z]$, it is possible. This is because the additional noise introduced by noise traders does not affect agents’ beliefs about the mean of z . Again, this highlights how direct observations of subjective forecasts can be helpful for disentangling

different hypotheses that are mathematically closely related.

One important caveat here (as in all the examples of this section) is that the subjective mean variable b is unrelated to ξ . In addition to incorporating noise traders, Hassan and Mertens also assume that agents are subject to belief distortions. Namely, that agents have imperfect information about z and use observations of x to compute their forecasts. This type of signal-extraction has the result of making b a function of ξ . In that case, it becomes more difficult to differentiate these hypotheses using forecasts alone, because even average forecasts would depend on the disturbance introduced by noise traders.

A.3 Belief Distortions from Noisy Signals

Instead of working directly with distorting martingales, another common approach to introduce belief distortions is to assume that agents receive noisy signals of fundamentals, which they use to generate forecasts according to Bayes' rule. This section considers two such signal extraction problems, and discusses the types of belief distortions they imply. Two differences relative to the empirical specification in section 1.3 are that signal extraction with Bayesian updating generates (i) distortions in conditional variances as well as conditional means, and (ii) (possibly dynamic) correlations between current fundamental innovations and innovations to beliefs.

A.3.1 Signals about the Future

As in section 1.3, a scalar fundamental process follows the law of motion

$$z_t = \rho_z z_{t-1} + e_t^z, \quad e_t^z \stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \sigma_z^2).$$

In addition to observing this process, an economic agent also receives a subjective, noisy signal about future fundamentals:

$$s_t = z_{t+1} + e_t^s, \quad e_t^s \stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \sigma_s^2).$$

The innovation e_t^s is the source of exogenous variation in the agent's beliefs. It is completely independent of fundamentals. The agent incorporates the information from this signal, and forecasts future fundamentals according to the conditional distribution

$$z_{t+1}|z_t, s_t \stackrel{\mathbb{P}}{\sim} \mathcal{N}((1-k)\rho_z z_t + k s_t, (1-k)\sigma_z^2),$$

where $k \equiv \sigma_z^2/(\sigma_z^2 + \sigma_s^2)$ controls how much weight is given to the signal. When the signal is completely uninformative ($\sigma_s \rightarrow \infty$), this weight is zero. When the signal is perfectly revealing ($\sigma_s \rightarrow 0$), this weight is one.

Define the distorting martingale M_t in the following way: $M_0 = 1$ and

$$M_{t+1} = M_t \frac{p(z_{t+1}|z_t, s_t)}{p(z_{t+1}|z_t)}.$$

Some algebraic manipulation shows that

$$M_{t+1} = M_t \left(\frac{\sigma_z}{\sigma} \right) \exp \left(\frac{1}{2} \left[\frac{1}{\sigma_z^2} - \frac{1}{\sigma^2} \right] e_{t+1}^z - \frac{1}{2} \frac{b_t^2}{\sigma^2} + \frac{b_t}{\sigma^2} e_{t+1}^z \right),$$

where $b_t \equiv k(s_t - \rho_z z_t)$ and $\sigma \equiv (1-k)\sigma_z^2$. Using this martingale to define a subjective probability measure \mathbb{Q} as in section 1.3, it follows that the evolution of fundamentals under both objective and subjective probabilities can be written in the form

$$\begin{aligned} z_t &= \rho_z z_{t-1} + e_t^z, \quad e_t^z \stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \sigma_z^2) \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(b_{t-1}, \sigma^2) \\ b_t &= \rho_b b_{t-1} + \tau e_{t+1}^z + e_t^b, \quad e_t^b \stackrel{\mathbb{P}, \mathbb{Q}}{\sim} \mathcal{N}(0, \sigma_b^2), \end{aligned}$$

with the restrictions $\sigma_b = k\sigma_s$, $\tau = k$, and $\rho_b = 0$. Note that there are two main differences relative to the dynamics of belief distortions in (1.5). First, there is a distortion to the conditional variance of z_t — by a constant factor $(1-k)$ — and second, the current innovation to z_t is (intertemporally) correlated with b_t . The reason for the first difference is that the signal provides additional information concerning future fundamentals, which reduces the subjective uncertainty of the agent in proportion to its precision. The reason for the second difference is that the signal is informative about future productivity; if the agent sees a high signal, he increases his conditional mean estimate of future productivity in proportion to the signal's precision.

A.3.2 Limited Information

In the previous subsection, the agent could observe current and past fundamentals, and received *additional* information about future fundamentals in the form of a noisy signal. In this subsection, I suppose instead that the agent has *less* information; he never observes the fundamental process and must base his forecasts exclusively on a sequence of noisy signals.

In each period, the agent observes a noisy signal of the current realization of z_t ,

$$s_t = z_t + e_t^s, \quad e_t^s \stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \sigma_s^2).$$

He incorporates the information from this signal (and all past signals) to forecast future fundamentals. Let $\{\mathfrak{F}_t^s\}_{t \geq 0}$ denote the natural filtration generated by these signals. Note that $\mathfrak{F}_t^s \subseteq \mathfrak{F}_t$, with equality only in the limiting case that the signal is perfectly informative ($\sigma_s = 0$). His forecasting distribution is

$$z_{t+1} | \mathfrak{F}_t^s \stackrel{\mathbb{P}}{\sim} \mathcal{N}(\rho_z \mu_t, \sigma^2),$$

where $\mu_t = (1 - k)\rho_z\mu_{t-1} + ks_t$, $k \equiv \sigma^2/(\sigma^2 + \sigma_s^2)$, and σ^2 is the solution to the equation¹

$$\sigma^2 = \left(\frac{\sigma_s^2}{\sigma^2 + \sigma_s^2} \right) \rho_z^2 \sigma^2 + \sigma_z^2.$$

Define the distorting martingale M_t in the following way: $M_0 = 1$ and

$$M_{t+1} = M_t \frac{p(z_{t+1}|\mathfrak{F}_t^s)}{p(z_{t+1}|\mathfrak{F}_t)}.$$

As before, some algebraic manipulation shows that

$$M_{t+1} = M_t \left(\frac{\sigma_z}{\sigma} \right) \exp \left(\frac{1}{2} \left[\frac{1}{\sigma_z^2} - \frac{1}{\sigma^2} \right] e_{t+1}^z - \frac{1}{2} \frac{b_t^2}{\sigma^2} + \frac{b_t}{\sigma^2} e_{t+1}^z \right),$$

where $b_t \equiv \rho_z(\mu_t - z_t)$. Using this martingale to define a subjective probability measure \mathbb{Q} as in section 1.3, it follows that the evolution of fundamentals under both objective and subjective probabilities can be written in the form

$$\begin{aligned} z_t &= \rho_z z_{t-1} + e_t^z, & e_t^z &\stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \sigma_z^2) \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(b_{t-1}, \sigma^2) \\ b_t &= \rho_b b_{t-1} + \tau e_t^z + e_t^b, & e_t^b &\stackrel{\mathbb{P}, \mathbb{Q}}{\sim} \mathcal{N}(0, \sigma_b^2), \end{aligned} \tag{A.1}$$

with the restrictions $\sigma_b = \rho_z k \sigma_s$, $\tau = -\rho_z(1 - k)$, and $\rho_b = \rho_z(1 - k)$. As before, the two main differences relative to the dynamics in (1.5) are that the conditional variance of z_t is distorted, and the current innovation to z_t is correlated with b_t .

There is one important caveat in this case, however. The martingale M_t captures distortions about *future* fundamentals, but not about current or past fundamentals. But, due to the hypothesis that the agent never observes the fundamental process, his beliefs about current and past fundamentals are distorted as well. That means, for example, that

$$E^{\mathbb{P}}[z_{t-j}|\mathfrak{F}_t^s] \neq z_{t-j} \quad \text{for any } j \geq 0.$$

¹Here I assume for simplicity that prediction is taking place under “steady-state” learning.

This presents a difficulty because it amounts to a violation of absolute continuity between subjective and objective beliefs. Because z_t is never observed, it is always viewed as a random variable by the agent. From the objective (full-information) perspective of the model, however, all randomness associated with z_t is resolved once date t is reached. For example, if the realization of z_t is z^* , the objective probability that z_t is not equal to z^* is zero for any $t \geq 0$, but the subjective probability is always non-zero.

Fortunately, in the case that only a finite number of lags of the fundamental process are relevant for the agent's decisions, it is still mechanically possible to proceed in the same spirit. For illustration, suppose it is necessary to determine the joint distribution of z_{t+1} and z_t . This is the most common case, but additional lags can be accommodated in the same way. Define the vector $x_t \equiv (z_t, z_{t-1})'$, and write the observation and state equations as

$$\begin{aligned} s_t &= Ax_t + De_t^s \\ x_t &= Bx_{t-1} + Ce_t^z, \end{aligned}$$

where $A \equiv (1, 0)$, $D \equiv (1, 0)'$, $C \equiv (1, 0)'$, and

$$B \equiv \begin{bmatrix} \rho_z & 0 \\ 1 & 0 \end{bmatrix}.$$

The agent's forecasting distribution is

$$x_{t+1} | \mathfrak{F}_t^s \stackrel{\mathbb{P}}{\sim} \mathcal{N}(B\mu_t, \Sigma),$$

where $\mu_t = (I_2 - KA)B\mu_{t-1} + Ks_t$, $K \equiv \Sigma A'(A\Sigma A' + \sigma_s^2 DD')^{-1}$, and Σ is the solution to the discrete algebraic Ricatti equation

$$\Sigma = B[I - \Sigma A'(A\Sigma A' + \sigma_s^2 DD')^{-1}A]\Sigma B' + \sigma_z^2 CC'.$$

As before, it is possible to write the subjective dynamics of x_t “as if” it were observable to the agent but subject to stochastic distortions given by $b_t \equiv B(\mu_t - x_t)$. Specifically, the dynamic system is of the form²

$$\begin{aligned} x_t &= Bx_{t-1} + e_t^x, & e_t^x &\overset{\mathbb{P}}{\sim} \mathcal{N}(0, \sigma_z^2 DD') \overset{\mathbb{Q}}{\sim} \mathcal{N}(b_{t-1}, \Sigma) \\ b_t &= Rb_{t-1} + Te_t^x + e_t^b, & e_t^b &\overset{\mathbb{P}, \mathbb{Q}}{\sim} \mathcal{N}(0, \Sigma_b), \end{aligned}$$

with the restrictions $\Sigma_b = BKD$, $T = -B(I - KA)$, and $R = B(I - KA)$. In this case, in addition to having distortions of the conditional covariance of fundamentals and correlation between fundamentals and the degree of distortions, the distortion process b_t is now a two-dimensional vector.

A.4 An Optimal Policy Example

This section presents a simple example to illustrate how the optimal policy response to advance information and distorted beliefs about future fundamentals are different. Specifically, I show that policy is more accomodative to exogenous innovations reflecting advance information about future cost-push shocks than it is to innovations that reflect distorted beliefs. This is because advance information generates adverse cost-push effects both through agents’ anticipation of future conditions, and through subsequent realizations of those conditions. Distortions in agents’ beliefs only engender a policy response through the first channel.

The example takes place in the context of the New Keynesian model presented in [Adam and Woodford \(2012\)](#). I will omit the details of the model and jump directly to the optimal monetary policy problem. The policymaker is charged with choosing maximizing the level

²Whenever the notation \mathcal{N} is used with a singular covariance matrix, what is meant is a “singular” normal distribution, in the sense of Anderson (1984).

of expected utility of a representative household,

$$E_0^{\mathbb{P}} \sum_{t=0}^{\infty} \beta^t U(Y_t, \Delta_t; \xi_t),$$

where Y_t is aggregate output, Δ_t is an index of price dispersion arising from staggered, monopolistic price-setting in product markets, and ξ_t is a vector of exogenous disturbances. His maximization is subject to the equilibrium conditions characterizing equilibrium in the private sector,

$$\begin{aligned} Z_t &= z(Y_t; \xi_t) + \alpha \beta E_t^{\mathbb{Q}}[\Phi(Z_{t+1})] \\ \Delta_t &= \tilde{h}(\Delta_{t-1}, K_t/F_t), \end{aligned}$$

together with a pre-commitment constraint necessary to deliver a time-invariant solution. Here the notation exactly follows [Adam and Woodford \(2012\)](#), with the exception that the probability measures \mathbb{P} and \mathbb{Q} are explicitly labeled. As in their paper, the objective is a paternalistic one. The policymaker understands that private sector beliefs may be different from his own, but still operates under the assumption that those beliefs are nevertheless computed according to a well-defined probability measure.

Let \tilde{M}_t denote a distorting martingale with unit expectation under \mathbb{Q} that governs the change of measure from \mathbb{P} to \mathbb{Q} . That is, for any $\mathcal{A} \in \mathfrak{F}_t$, the following relation holds:

$$\mathbb{P}(\mathcal{A}) = E^{\mathbb{Q}}[1_{\mathcal{A}} \tilde{M}_t]$$

Using this martingale process, the Lagrangian for the policy problem can be written under

the distorted probability measure:

$$E_0^{\mathbb{Q}} \sum_{t=0}^{\infty} \beta^t \tilde{M}_t \left\{ U(Y_t, \Delta_t; \xi_t) + \gamma_t [\tilde{h}(\Delta_{t-1}, K_t/F_t) - \Delta_t] + \Gamma'_t [z(Y_t; \xi_t) + \alpha \beta \Phi(Z_{t+1}) - Z_t] \right\} \\ + \alpha \Gamma'_{-1} \Phi(Z_0),$$

The associated first order conditions are:

$$U_t(Y_t, \Delta_t; \xi_t) + \Gamma'_t z_Y(Y_t; \xi_t) = 0 \quad (\text{A.2})$$

$$-\gamma_t \tilde{h}_2(\Delta_{t-1}, K_t/F_t) \frac{K_t}{F_t^2} - \Gamma_{1t} + \alpha \Gamma'_{t-1} D_1(K_t/F_t) = 0 \quad (\text{A.3})$$

$$\gamma_t \tilde{h}_2(\Delta_{t-1}, K_t/F_t) \frac{1}{F_t} - \Gamma_{2t} + \alpha \Gamma'_{t-1} D_2(K_t/F_t) = 0 \quad (\text{A.4})$$

$$U_{\Delta}(Y_t, \Delta_t; \xi_t) - \gamma_t + \beta E_t^{\mathbb{Q}}[(\tilde{M}_{t+1}/\tilde{M}_t) \gamma_{t+1} \tilde{h}_1(\Delta_t, K_{t+1}/F_{t+1})] = 0 \quad (\text{A.5})$$

Computing a log-linear approximation of these conditions around the deterministic steady-state (still under \mathbb{Q}), optimal policy ensures that the following system of equations in inflation (π_t) and the output gap (x_t) is always satisfied:

$$\pi_t = \kappa x_t + \beta E_t^{\mathbb{Q}}[\pi_{t+1}] + u_t \quad (\text{A.6})$$

$$\xi_{\pi} \pi_t + \lambda_x (x_t - x_{t-1}) = 0.$$

In the first equation, u_t is a composite cost-push disturbance defined as a function of the exogenous fundamentals in the economy. The parameters $\kappa, \xi_{\pi}, \lambda_x > 0$ are reduced-form parameters that are functions of the underlying deep parameters of the model. The only differences in this example from [Adam and Woodford \(2012\)](#) are that the log-linear approximation has been taken under \mathbb{Q} rather than \mathbb{P} , and the policymaker does not have a concern for robustness against the possibility that private sector beliefs can change in ways that represent *worst-case* outcomes from the perspective of his policy objective.

Suppose that the exogenous process u_t follows dynamics like those in system (1.5), but

for the purpose of this example, without any persistence:

$$\begin{aligned} u_t &= a_{t-1} + e_t^u, & e_t^u &\stackrel{\mathbb{P}}{\sim} \mathcal{N}(0, \sigma_u^2) \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(b_{t-1}, \sigma_u^2) \\ a_t &= e_t^a, & e_t^a &\stackrel{\mathbb{P}, \mathbb{Q}}{\sim} \mathcal{N}(0, \sigma_a^2) \\ b_t &= e_t^b, & e_t^b &\stackrel{\mathbb{P}, \mathbb{Q}}{\sim} \mathcal{N}(0, \sigma_b^2) \end{aligned}$$

Solving system (A.6) using these dynamics for u_t under \mathbb{Q} , the optimal evolution of inflation and the output gap are

$$\begin{aligned} \pi_t &= \gamma u_t + \beta \gamma^2 (a_t + b_t) + \frac{\lambda_x}{\xi_\pi} (1 - \gamma) x_{t-1} \\ x_t &= -\frac{\xi_\pi}{\lambda_x} \gamma u_t - \beta \frac{\xi_\pi}{\lambda_x} \gamma^2 (a_t + b_t) + \gamma x_{t-1}, \end{aligned} \tag{A.7}$$

where $0, \gamma < 1$ is the smaller of the two real roots of

$$\gamma^2 - \left(1 + \frac{1}{\beta} + \frac{\kappa \xi_\pi}{\lambda_x \beta}\right) \gamma + \frac{1}{\beta} = 0.$$

Let $R_t[\pi_{t+\tau} | e_t = e^*] \equiv E_{t-1}^{\mathbb{P}}[\pi_{t+\tau} | e_t = e^*] - E_{t-1}^{\mathbb{P}}[\pi_{t+\tau}]$ denote the impulse response function of inflation in period $\tau \geq 0$ to an innovation $e_t = e^*$ in period t . Now consider the difference between the response of inflation to an exogenous unit innovation in e_t^a and e_t^b . The relations in (A.7) and the exogenous dynamics of u_t imply that

$$R_t[\pi_{t+\tau} | e_t^a = 1] - R_t[\pi_{t+\tau} | e_t^b = 1] = \gamma 1_{\{\tau=0\}} + \gamma^\tau \frac{\lambda_x}{\xi_\pi} (1 - \gamma) 1_{\{\tau>0\}}.$$

Because $\lambda, \xi_\pi > 0$ and $0 < \gamma < 1$, this difference is strictly positive for all horizons $\tau > 0$. This shows that it is optimal for the policymaker to allow inflation to respond more to innovations that reflect advance information than to innovations that reflect distortions in agents' beliefs. The same logic shows that the output gap x_t exhibits a greater negative response to advance information compared to distorted beliefs.

Appendix B

Appendix for Chapter 2

B.1 Robustness

Our baseline estimate of macro uncertainty $\bar{\mathcal{U}}_t^y(h)$ is constructed as the cross-sectional average of the individual uncertainties $\mathcal{U}_{jt}^y(h)$, and each of these is based on evaluating (2.11) at the posterior mean, over the full sample, of the state and parameters of the stochastic volatility model (i.e., $\{\log(\sigma_{jt}^y)^2\}, \alpha_j^y, \beta_j^y$, and τ_j^y) and the OLS parameter estimates from the forecasting model (i.e., $\phi_j^y, \gamma_j^F(L)$, and $\gamma_j^W(L)$). This section assesses robustness of the results to these assumptions.

B.1.1 Macro Uncertainty Factor

We first entertain the possibility that uncertainty has a factor structure. In such a case, macro uncertainty at each t is a vector given by the common factor $F_t^{\mathcal{U}}(h)$ in

$$\log \mathcal{U}_{jt}^y(h) = c_j^{\mathcal{U}}(h) + \Lambda_{hj}^{\mathcal{U}} F_t^{\mathcal{U}}(h) + e_{jt}^{\mathcal{U}}(h). \quad (\text{B.1})$$

Macro uncertainty is then summarized by $F_t^{\mathcal{U}}(h)$ while idiosyncratic uncertainty is $e_{hjt}^{\mathcal{U}}$. Although $\mathcal{U}_{jt}^y(h)$ is always positive, the principal components estimates do not constrain the

(normalized) estimated factors themselves to be positive. The log specification is therefore used to insure that both the domain and the range of the function (B.1) take on values on the entire real line \mathbb{R} . As a consequence of this log specification, our PCA estimate of macro uncertainty $\mathcal{U}_t^y(h)$ is the exponential of the PCA estimate $\hat{F}_t^{\mathcal{U}}(h)$. Let $\hat{\mathcal{U}}_t^y(h) \equiv \exp\left(\hat{F}_t^{\mathcal{U}}(h)\right)$. To obtain such an estimate, we first need an estimate of the common (log) uncertainty factor $F_t^{\mathcal{U}}(h)$. As many uncertainty series appear non-stationary, this estimate is defined by $\hat{F}_t^{\mathcal{U}}(h) = \sum_{k=2}^t \hat{f}_k^{\mathcal{U}}(h)$, where $f_t^{\mathcal{U}}(h)$ is an $r_{\mathcal{U}} \times 1$ vector comprised of the $r_{\mathcal{U}}$ principal components of $\Delta \log \mathcal{U}_{jt}^y(h)$.¹ As discussed in Bai and Ng (2004), this differencing-recumulating approach ensures that the factors are consistently estimated when the idiosyncratic errors are potentially non-stationary. Because of the differencing, the initial value in the sample of the common uncertainty factor, $\hat{F}_1^{\mathcal{U}}(h)$, is not identified. We initialize $\hat{F}_1^{\mathcal{U}}(h)$ to the average level of (log) uncertainty across all N series; mathematically, $\frac{1}{N} \sum_{j=1}^N \log \mathcal{U}_{j1}^y(h)$.

The problem of determining $r_{\mathcal{U}}$, the number of common uncertainty factors $f^{\mathcal{U}}(h)$, is non-standard because the individual uncertainty measures are themselves estimated. Existing criteria for determining the number of factors do not take the first step estimation error into account and will likely overestimate the number of factors. However, there is strong evidence of a factor structure as the largest eigenvalue of forecast error variance is distinctly large. In particular, the first principal component of $\mathcal{U}_{jt}^y(h)$ explains 11% of the variance of the forecast errors for $h = 1$, 14% for $h = 3$, and 22% for $h = 12$. We take $r_{\mathcal{U}}$ to be one, which facilitates comparison with the base-case estimate $\bar{\mathcal{U}}_t^y(h)$ that is based on simple averaging. We also calibrate the uncertainty factor $\hat{\mathcal{U}}_t^y(h)$ to have the same mean and standard deviation as $\bar{\mathcal{U}}_t(h)$ over the sample.

The right panel of Table 2.1 shows that the results using $\hat{\mathcal{U}}_t^y(h)$ are qualitatively and quantitatively similar to the base-case. The relative importance of the uncertainty factor

¹We observe $\log \mathcal{U}_{jt}^y(h)$, a data matrix with T time-series observations and N cross-section observations. The first differenced data yield a $(T-1) \times N$ vector of stationary variables. Let $f^{\mathcal{U}}(h) \equiv (f_1^{\mathcal{U}}(h), f_2^{\mathcal{U}}(h), \dots, f_T^{\mathcal{U}}(h))$ and $\Lambda_{\mathcal{U}} = (\Lambda_{\mathcal{U}\infty}, \Lambda_{\mathcal{U}\in}, \dots, \Lambda_{\mathcal{U}N})'$. The principal component estimator of $f^{\mathcal{U}}(h)$ is the $T-1$ times the $r_{\mathcal{U}}$ eigenvectors corresponding to the first $r_{\mathcal{U}}$ largest eigenvalues of the $(T-1)(T-1)$ matrix $(\Delta \log \mathcal{U}_{jt}^y(h)) (\Delta \log \mathcal{U}_{jt}^y(h))'$.

and idiosyncratic uncertainty is summarized in a $R_{jt}^2(h)$ statistic analogous to (2.14). The main finding continues to be that variations in macro uncertainty constitute a larger fraction of variations in individual uncertainty measures at longer horizons, and during recessions.

B.1.2 Alternative Estimates of Uncertainty

We next consider alternative estimates of individual uncertainty, and alternative ways of aggregating these estimates to get macro uncertainty. The base-case implementation only requires one evaluation of uncertainty for each series j since the posterior mean of each parameter is one dimensional. Specifically, for $h = 1$, uncertainty in the variable j evaluated at the s th Monte Carlo draw is

$$\mathcal{U}_{jst}(h)(\theta_{js}, x_{jst}) = \sqrt{\exp(\alpha_{js} + \tau_{js}^2/2 + \beta_{js}x_{jst})},$$

where $x_{jst} \equiv \ln(\sigma_{jst}^y)^2$. When the function above is evaluated at the posterior mean (over all $s = 1, \dots, S$ draws) of the parameters, we denote that $\mathcal{U}_{jt}(h)(\bar{\theta}_j, \bar{x}_{jt})$. In this notation, our base case uncertainty estimate for the series j is $\mathcal{U}_{jt}(h)(\bar{\theta}_j, \bar{x}_{jt})$. But an uncertainty estimate can also be obtained for each draw of the hyperparameters in the model for series j . Thus one can also estimate $\mathcal{U}_{jt}^y(h)$ by the posterior mean of the draws of uncertainty for series j . In this case we define individual uncertainty as $\mathcal{U}_{jt}^S(h) = \frac{1}{S} \sum_{s=1}^S \mathcal{U}_{jst}(h)(\theta_{js}, x_{jst})$, where the superscript S denotes all S draws are used in the computation.² Instead of the posterior mean, it is also possible to consider other location statistics. Let $\mathcal{U}_{jt}^{[s]}(h)$ be the s -th percentile draw in the sorted sequence of $\{\mathcal{U}_{jst}(h)\}_{s=1}^S$. If $[s]$ is 50, the median obtains. We use the 90th and the 10th percentiles of the posterior distribution of $\mathcal{U}_{jst}(h)(\theta_{js}, x_{jst})$ to assess how extreme values of individual uncertainty affect aggregate uncertainty. These are denoted $\bar{\mathcal{U}}_t^{10}(h)$ and $\bar{\mathcal{U}}_t^{90}(h)$, respectively. Since we have three ways of estimating individ-

²To estimate the latter requires saving every posterior draw of $\mathcal{U}_{jst}(h)(\theta_{js}, x_{jst})$ and is considerably more computationally demanding than the base-case where uncertainty is evaluated once at the mean of the parameters.

$\mathcal{U}_t(h)$	Aggregator	$\mathcal{U}_{jt}(h)$
Baseline CSA: $\overline{\mathcal{U}}_t(h)$	CSA	$\mathcal{U}_{jt}(h) (\overline{\theta}_j, \overline{x}_{jt})$
Baseline PCA: $\widehat{\mathcal{U}}_t(h)$	PCA	$\mathcal{U}_{jt}(h) (\overline{\theta}_j, \overline{x}_{jt})$
Posterior Mean CSA: $\overline{\mathcal{U}}_t^S(h)$	CSA	$\frac{1}{S} \sum_{s=1}^S \mathcal{U}_{jst}(h) (\theta_{js}, x_{jst})$
Posterior Mean PCA: $\widehat{\mathcal{U}}_t^S(h)$	PCA	$\frac{1}{S} \sum_{s=1}^S \mathcal{U}_{jst}(h) (\theta_{js}, x_{jst})$
Posterior s -Percentile CSA: $\overline{\mathcal{U}}_t^{[s]}(h)$	CSA	$\mathcal{U}_{jt}^{[s]}(h)$
Posterior s -Percentile PCA: $\widehat{\mathcal{U}}_t^{[s]}(h)$	PCA	$\mathcal{U}_{jt}^{[s]}(h)$

Table B.1: Different Aggregation Methods.

ual uncertainties two ways of aggregating them, we have six measures of macro uncertainty summarized in Table B.1. In that table, CSA stands for simple averaging over N_y series, and PCA stands for for the principal component of the N_y individual uncertainties constructed using the methodology as discussed above.

Figure (B.1) shows the baseline and posterior mean estimates of aggregate uncertainty when $h = 1$. Each of these measures are highly correlated with one another. Indeed, the estimates based on the average across draws of the parameters versus the posterior mean of the uncertainty draws are virtually indistinguishable. The estimates based on cross-section averaging are also very highly correlated with those based on the principal component estimates. Given the similarity between the CSA and PCA estimates, Figure (B.2) shows our base-case estimate of uncertainty $\overline{\mathcal{U}}_t(h)$, the CSA variant of $\overline{\mathcal{U}}_t^S(h)$, along with the CSA variant of $\overline{\mathcal{U}}_t^{10}(h)$ and $\overline{\mathcal{U}}_t^{90}(h)$. As for the above variations, different percentiles of the distribution have the effect of shifting our estimate of uncertainty by a constant amount only but do not much affect the dynamics of our uncertainty estimates. The 90th and 10th percentiles of the distribution have a correlation with our baseline estimate each in excess of 0.998. We conclude that results regarding the number of large uncertainty episodes, their timing, or their dynamic relation with economic activity are robust to using more extreme estimates of individual uncertainty. Overall, the results suggest that the findings reported above are not sensitive to using these alterative estimates of aggregate uncertainty.

Finally, we consider using GARCH or EGARCH to estimate the volatility of individual




Figure B.1: Different estimates of macro uncertainty when $h = 1$. Baseline CSA is $\bar{U}_t(1) = \frac{1}{N_y} \sum_{j=1}^{N_y} \mathcal{U}_{jt}(1) (\bar{\theta}_j, \bar{x}_{jt})$. Baseline PCA shows the principal component based on $\mathcal{U}_{jt}(1) (\bar{\theta}_j, \bar{x}_{jt})$. Posterior mean CSA is the cross-section average of $\frac{1}{S} \sum_{s=1}^S \mathcal{U}_{jst}(1) (\theta_{js}, x_{jst})$. Posterior mean PCA shows the first principal component based on $\frac{1}{S} \sum_{s=1}^S \mathcal{U}_{jst}(1) (\theta_{js}, x_{jst})$. The full sample spans the period 1960:01-2011:12.




Figure B.2: Percentile-based estimates of aggregate uncertainty when $h = 1$. Baseline denotes our base-case CSA estimate of macro uncertainty: $\bar{u}_t(1) = \frac{1}{N_y} \sum_{j=1}^{N_y} \mathcal{U}_{jt}(1) (\bar{\theta}_j, \bar{x}_{jt})$ and $\bar{\theta}_j$ and \bar{x}_{jt} are posterior means over S draws. Posterior mean CSA is $\bar{u}_t(1) = \frac{1}{N_y} \sum_{j=1}^{N_y} \frac{1}{S} \sum_{s=1}^S \mathcal{U}_{jst}(1) (\theta_{js}, x_{jst})$. The posterior percentile- s CSA is $\bar{u}_t(1) = \frac{1}{N_y} \sum_{j=1}^{N_y} \mathcal{U}_{jt}^{[s]}(1)$ where $\mathcal{U}_{jt}^{[s]}(1)$ is the s -th percentile draw in the ordered sequence of $\mathcal{U}_{jst}(1)(\theta_{js}, x_{jst})$, for $s = 1, \dots, S$. The sample spans the period 1960:01-2011:12.

series. Figure (B.3) shows that, when we aggregate in exactly the same way, our estimates of aggregate uncertainty over time are very similar to the baseline stochastic volatility case. Results based on the GARCH/EGARCH estimates indicate the number and timing of big uncertainty episodes, as well as the persistence of uncertainty, is very similar to that reported here using our base-case measure of macro uncertainty. What is different is the real effect of uncertainty innovations from a VAR, once orthogonalized shocks are analyzed. This is to be expected because GARCH type models (unlike stochastic volatility) have a shock to the second moment that is not independent of the first moment. This is inconsistent with the assumptions of an independent uncertainty shock presumed in the uncertainty literature. Using a GARCH-based uncertainty index thus creates additional identification problems that are beyond the scope of this paper.

B.1.3 Recursive Out-of-Sample Estimation

We next consider the sensitivity of the forecasting parameters ϕ_j^y , $\gamma_j^F(L)$, and $\gamma_j^W(L)$ to the estimation sample. Instead of full sample estimation (and hence in-sample forecasts), we also form out-of-sample forecasts for the monthly macro dataset.³ This procedure involves fully recursive factor estimation and parameter estimation using data only through time t for forecasting at time $t+1$. Notice that, since the forecasting parameters evolve over time as new data becomes available, such recursive forecasts are informative about the extent to which parameter instability in the conditional mean forecasting relation influences the uncertainty estimates. We use the first 10 years of data ($t = 1, 2, \dots, 120$, 1959:01-1969:01) as an initial estimation period to estimate both the factors and the parameters of the conditional mean (forecasting) regression, and to perform model selection. Next, the forecasting regressions are run over the period $t = 1959:01, \dots, 1969:01$, and the values of the regressors at $t = 1969:01$ are used to forecast $y_{1969:02}$. All parameters, factors and model selection criteria are then re-estimated from 1959:01 through 1969:02, and forecasts are recomputed for $y_{1969:03}$, and

³This procedure closely follows the real-time simulation procedure of [Stock and Watson \(2002a\)](#).




Figure B.3: EGARCH Aggregate Uncertainty: $\bar{U}_t^y(1)$ computed using baseline stochastic volatility estimates, and EGARCH(1,1) estimates with t -distributed errors. Aggregate uncertainty is calculated as before, using a simple cross-sectional average. The data are monthly and span the period 1960:07-2011:12.

so on, until the final out-of-sample forecast is made for $y_{j2011:12}$. Since our dataset has 622 months total, this leaves $502=622-120$ forecast errors. The forecast error variances are used to compute $\overline{U}_{jt}^y(h)$, and averaging over j gives macro uncertainty. The resulting uncertainty estimate is plotted in Figure B.4 along with the original estimate. The measure is extremely highly correlated with that based on in-sample forecasts.⁴ Although use of the full sample slightly under-states the level of uncertainty, it does an excellent job of capturing its time-series variation, only influencing the estimates by a constant amount. We can confirm that our VAR analysis is little affected by whether we use out-of-sample or in-sample forecasts, having virtually no bearing on the number of uncertainty episodes, their timing, or their dynamic relationship with economic activity. These findings are consistent with evidence that dynamic factor analysis provides robustness against the temporal parameter instability that often plagues low-dimensional forecasting regressions (Stock and Watson, 2002b). The reason is that such instabilities can “average out” in the construction of common factors if the instability is sufficiently dissimilar from one series to the next. In the recursive VAR estimation the parameters of the forecasting relation change every period, so this speaks directly to the question of the role played by parameter stability in our estimates.

The recursive out-of-sample approach is only feasible in the $h = 1$ case. This is because we obtain our estimates of uncertainty for $h > 1$ by are computed once by rolling ahead one-step-ahead forecasts from the VAR stacked in companion form. By design, this approach relies on the parameters of the VAR being fixed over the sample. Nevertheless we find the robustness of the results in the $h = 1$ case along this dimension to be comforting and suggestive of what would be likely for the other cases.

⁴Note that this measure is feasible to compute only for $h = 1$. The multi-step ahead forecasts that are needed for uncertainty with $h > 1$ are computed once by rolling forward one-step ahead forecasts from the VAR. Recomputing the VAR in every time period would require recomputing uncertainty in every time period, which is not possible in reasonable time.




Figure B.4: Uncertainty factor based on recursive forecasts. This plot displays $\bar{U}_t^y(h)$ based on forecasts which use information from the full sample (“Baseline”), and based on recursively computed out-of-sample forecasts (“Real-time”), expressed in standardized units. The recursive forecasting procedure involves estimating model parameters and predictor variables only using information available up to time t . A training sample of 10 years (120 observations) is used to compute the first out-of-sample forecast, for 1970:01. The full sample spans the period 1960:01–2011:12.

B.1.4 Bloom (2009) VAR

The Bloom VAR results thus far have used an ordering that puts uncertainty second in a list of eight variables, following [Bloom \(2009\)](#). Table [B.2](#) reports VAR variance decomposition results with uncertainty ordered last to allow uncertainty to respond contemporaneously to the five variables ordered after it. Figure [B.5](#) reports the impulse responses to orthogonal shocks created from a Cholesky decomposition of the VAR with this alternative ordering. Some variations previously attributed to uncertainty are now allocated to the orthogonalized innovations in the fed funds rate, wages, CPI, hours, employment, and industrial production. This is not surprising because our measure of uncertainty is contemporaneously correlated with these measures of economic activity, thus once we remove the variation in uncertainty that is attributable to these correlations, the effect is smaller. We again caution, however, that these results as well as the previous ones tell us only about dynamic correlations (not true causality) and differ only because of a change in the assumption about the timing of shocks. For the sake of comparison, the last column of Table [B.2](#) reports results with VXO ordered last. As documented earlier, stock market volatility and uncertainty are correlated but have significant independent variations. As expected, because our measures of uncertainty are more highly contemporaneously correlated with real activity than is VXO, the effect on production, employment, and hours attributed to uncertainty shocks is smaller compared to the results in Table 2 when uncertainty is ordered second. By contrast, the decomposition of forecast error variances to VXO shocks is not greatly affected by the ordering of VXO in the VAR, implying that VXO shocks are not as strongly contemporaneously correlated with the five real activity variables in the system as are our uncertainty estimates. These results reinforce the conclusion that the stock market can move significantly in the absence changes in fundamentals in the economy. It is thus not a good proxy for macroeconomic uncertainty, which we have found does move with these fundamentals.

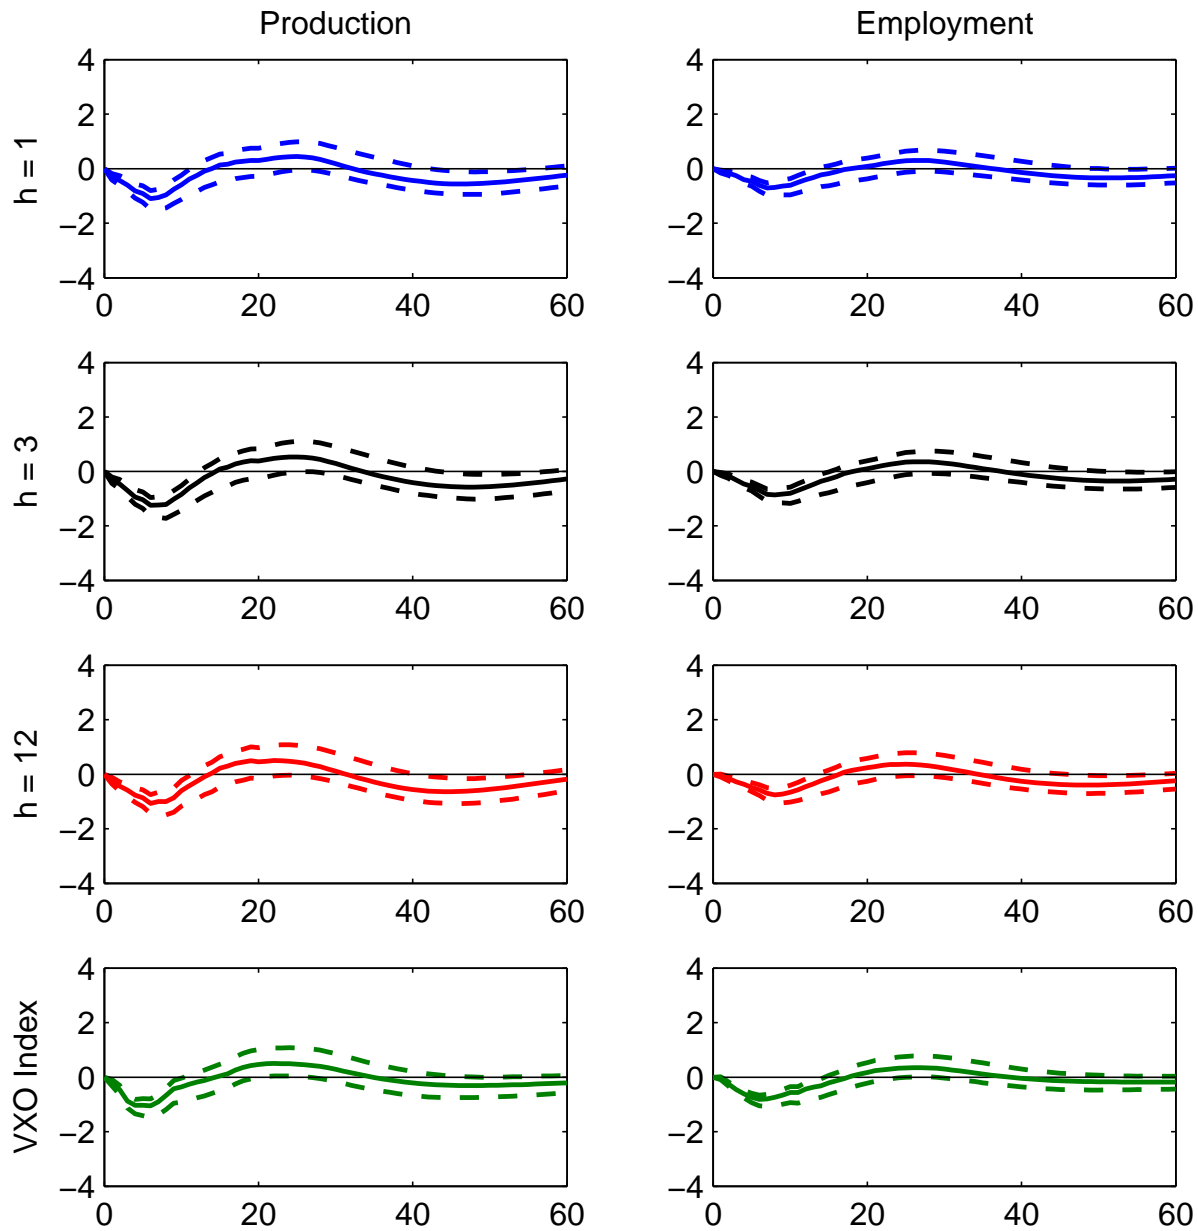


Figure B.5: Eight-variable VAR(12) with uncertainty ordered last. Uncertainty is measured using the VXO Index or $\bar{U}_t^y(h)$ for $h = 1, 3, 12$ as a measure of uncertainty. Each VAR(12) contains, in the following order: log(S&P 500 Index), federal funds rate, log(wages), log(CPI), hours, log(employment), log(industrial production), and *uncertainty*. All shocks are a 4 standard deviation impulse, which is the same magnitude considered in Bloom (2009) Figure A.1. As in Bloom (2009), all variables are HP filtered, except for the uncertainty measures, which enter in raw levels. The data are monthly and span the period 1960:07-2011:12.

Variance Decompositions from VAR(12)
(Uncertainty Ordered Last)

k	$\overline{U}(1)$	$\overline{U}(3)$	$\overline{U}(12)$	VXO
Production				
1	0.00	0.00	0.00	0.00
3	1.16	1.31	1.03	1.04
12	6.18	8.95	6.11	5.84
∞	5.51	7.26	6.33	4.14
max	6.78	9.45	6.62	7.19
max k	10	10	10	8
Employment				
1	0.00	0.00	0.00	0.00
3	0.60	0.59	0.43	1.11
12	5.97	9.20	6.58	8.88
∞	4.99	7.03	6.18	5.18
max	6.05	9.20	6.58	9.61
max k	11	12	12	9
Hours				
1	0.00	0.00	0.00	0.00
3	1.42	1.57	0.89	1.70
12	5.82	8.00	5.56	7.12
∞	5.94	7.97	6.81	5.98
max	6.21	8.40	6.81	7.86
max k	8	10	38	8

Table B.2: Eight-variable VAR(12) using the VXO Index or $\overline{U}_t^y(h)$ for $h = 1, 3, 12$ as a measure of uncertainty, estimated from the monthly macro dataset. Each VAR(12) contains, in the following order: log(S&P 500 Index), federal funds rate, log(wages), log(CPI), hours, log(employment), log(industrial production), and *uncertainty*. As in Bloom (2009), all variables are HP filtered, except for the uncertainty measures, which enter in raw levels. The data are monthly and span the period 1960:07-2011:12.

B.2 Data Appendix

The first dataset, denoted X^m , is an updated version of the of the 132 mostly macroeconomic series used in [Ludvigson and Ng \(2010\)](#). The 132 macro series in X^m are selected to represent broad categories of macroeconomic time series: real output and income, employment and hours, real retail, manufacturing and trade sales, consumer spending, housing starts, inventories and inventory sales ratios, orders and unfilled orders, compensation and labor costs, capacity utilization measures, price indexes, bond and stock market indexes, and foreign exchange measures.

The 147 financial series in X^f consists of a number of indicators measuring the behavior of a broad cross-section of asset returns, as well as some aggregate financial indicators not included in the macro dataset. These data include valuation ratios such as the dividend-price ratio and earnings-price ratio, growth rates of aggregate dividends and prices, default and term spreads, yields on corporate bonds of different ratings grades, yields on Treasuries and yield spreads, and a broad cross-section of industry equity returns. Following [Fama and French \(1992\)](#), returns on 100 portfolios of equities sorted into 10 size and 10 book-market categories. The dataset X^f also includes a group of variables we call “risk-factors,” since they have been used in cross-sectional or time-series studies to uncover variation in the market risk-premium. These risk-factors include the three [Fama and French \(1993\)](#) risk factors, namely the excess return on the market MKT_t , the “small-minus-big” (SMB_t) and “high-minus-low” (HML_t) portfolio returns, the momentum factor UMD_t , the bond risk premia factor of [Cochrane \(2005\)](#), and the small stock value spread $R15 - R11$.

The raw data used to form factors are always transformed to achieve stationarity. In addition, when forming forecasting factors from the large macro and financial datasets, the raw data (which are in different units) are standardized before performing PCA. When forming common uncertainty from estimates of individual uncertainty, the raw data (which are in this case in the same units) are demeaned, but we do not divide by the observation’s standard deviation before performing PCA.

Throughout, the factors are estimated by the method of static principal components (PCA). Specifically, the $T \times r_F$ matrix \hat{F}_t is \sqrt{T} times the r_F eigenvectors corresponding to the r_F largest eigenvalues of the $T \times T$ matrix $xx'/(TN)$ in decreasing order. In large samples (when $\sqrt{T}/N \rightarrow \infty$), [Bai and Ng \(2006\)](#) show that the estimates \hat{F}_t can be treated as though they were observed in the subsequent forecasting regression. There is no need to correct standard errors for uncertainty in this estimate, unlike the generated regressor case analyzed in [Pagan \(1984\)](#) when N is fixed. This asymptotic result allows for time variation in the volatility of the forecast error.

B.2.1 Macro Dataset

This appendix lists the short name of each series in the macro dataset, its code in the source database, the transformation applied to the series, and a brief data description. All series are from the IHS Global Insights database, unless the source is listed (in parentheses) as FRED (St. Louis Federal Reserve Economic Data), BLS (Bureau of Labor Statistics), S (R. J. Shiller website), BEA (Bureau of Economic Analysis), IMF (IMF International Financial Statistics database), B (R Barnichon website), UM (Thomson Reuters/University of Michigan Surveys of Consumers) or AC (author's calculation). The data are available from 1959:01-2011:12.

Let X_{it} denote variable i observed at time t after e.g., logarithm and differencing transformation, and let X_{it}^A be the actual (untransformed) series. Let $\Delta = (1 - L)$ with $LX_{it} = X_{it-1}$. There are six possible transformations with the following codes:

- (1) Code lv : $X_{it} = X_{it}^A$.
- (2) Code Δlv : $X_{it} = X_{it}^A - X_{it-1}^A$.
- (3) Code $\Delta^2 lv$: $X_{it} = \Delta^2 X_{it}^A$.
- (4) Code ln : $X_{it} = \ln(X_{it}^A)$.
- (5) Code Δln : $X_{it} = \ln(X_{it}^A) - \ln(X_{it-1}^A)$.
- (6) Code $\Delta^2 ln$: $X_{it} = \Delta^2 \ln X_{it}^A$.

No.	Gp	Short Name	Code	Tran	Description
1	1	PI	M_14386177	Δln	Personal Income
6	1	IP: total	M_116460980	Δln	Industrial Production Index - Total Index
7	1	IP: products	M_116460981	Δln	Industrial Production Index - Products, Total
8	1	IP: final prod	M_116461268	Δln	Industrial Production Index - Final Products
9	1	IP: cons gds	M_116460982	Δln	Industrial Production Index - Consumer Goods
10	1	IP: cons dble	M_116460983	Δln	Industrial Production Index - Durable Consumer Goods
11	1	IP: cons nondble	M_116460988	Δln	Industrial Production Index - Nondurable Consumer Goods
12	1	IP: bus eqpt	M_116460995	Δln	Industrial Production Index - Business Equipment
13	1	IP: matls	M_116461002	Δln	Industrial Production Index - Materials
14	1	IP: dble matls	M_116461004	Δln	Industrial Production Index - Durable Goods Materials
15	1	IP: nondble matls	M_116461008	Δln	Industrial Production Index - Nondurable Goods Materials
16	1	IP: mfg	M_116461013	Δln	Industrial Production Index - Manufacturing
17	1	IP: res util	M_116461276	Δln	Industrial Production Index - Residential Utilities
18	1	IP: fuels	M_116461275	Δln	Industrial Production Index - Fuels
19	1	NAPM prodn	M_110157212	lv	Napm Production Index
20	1	Cap util	M_116461602	Δlv	Capacity Utilization

Table B.3: Output and Income

No.	Gp	Short Name	Code	Tran	Descripton
21	2	Help wanted indx	-	Δlv	Index Of Help-Wanted Advertising (B)
22	2	Help wanted/unemp	M.110156531	Δlv	Ratio of Help-Wanted Ads/No. Unemployed (AC)
23	2	Emp CPS total	M.110156467	Δln	Civilian Labor Force: Employed, Total
24	2	Emp CPS nonag	M.110156498	Δln	Civilian Labor Force: Employed, Nonagric.Industries
25	2	U: all	M.110156541	Δlv	Unemployment Rate: All Workers, 16 Years & Over
26	2	U: mean duration	M.110156528	Δlv	Unemp By Duration: Average Duration In Weeks
27	2	U ; 5 wks	M.110156527	Δln	Unemp. By Duration: Persons Unempl Less Than 5 Wks
28	2	U 5-14 wks	M.110156523	Δln	Unemp. By Duration: Persons Unempl 5 To 14 Wks
29	2	U 15+ wks	M.110156524	Δln	Unemp. By Duration: Persons Unempl 15 Wks +
30	2	U 15-26 wks	M.110156525	Δln	Unemp. By Duration: Persons Unempl 15 To 26 Wks
31	2	U 27+ wks	M.110156526	Δln	Unemp. By Duration: Persons Unempl 27 Wks +
32	2	UI claims	M.15186204	Δln	Initial Claims for Unemployment Insurance
33	2	Emp: total	M.123109146	Δln	Emp. Nonfarm Payrolls: Total Private
34	2	Emp: gds prod	M.123109172	Δln	Emp. Nonfarm Payrolls: Goods-Producing
35	2	Emp: mining	M.123109244	Δln	Emp. Nonfarm Payrolls: Mining
36	2	Emp: const	M.123109331	Δln	Emp. Nonfarm Payrolls: Construction
37	2	Emp: mfg	M.123109542	Δln	Emp. Nonfarm Payrolls: Manufacturing
38	2	Emp: dble gds	M.123109573	Δln	Emp. Nonfarm Payrolls: Durable Goods
39	2	Emp: nondbles	M.123110741	Δln	Emp. Nonfarm Payrolls: Nondurable Goods
40	2	Emp: services	M.123109193	Δln	Emp. Nonfarm Payrolls: Service-Providing
41	2	Emp: TTU	M.123111543	Δln	Emp. Nonfarm Payrolls: Trade, Transport., Utilities
42	2	Emp: wholesale	M.123111563	Δln	Emp. Nonfarm Payrolls: Wholesale Trade
43	2	Emp: retail	M.123111867	Δln	Emp. Nonfarm Payrolls: Retail Trade
44	2	Emp: FIRE	M.123112777	Δln	Emp. Nonfarm Payrolls: Financial Activities
45	2	Emp: Govt	M.123114411	Δln	Emp. Nonfarm Payrolls: Government
*46	2	Agg wkly hours	-	Δlv	Index of Aggregate Weekly Hours (BLS)
*47	2	Avg hrs	M.140687274	Δlv	Avg Weekly Hrs, Prod/Nonsup: Goods-Producing
*48	2	Overtime: mfg	M.123109554	Δlv	Avg Weekly Hrs, Prod/Nonsup: Mfg Overtime
*49	2	Avg hrs: mfg	M.14386098	Δlv	Average Weekly Hours, Mfg.
50	2	NAPM empl	M.110157206	lv	NAPM Employment Index
129	2	AHE: goods	M.123109182	$\Delta^2 ln$	Avg Hourly Earnings, Prod/Nonsup: Goods-Producing
130	2	AHE: const	M.123109341	$\Delta^2 ln$	Avg Hourly Earnings, Prod/Nonsup: Construction
131	2	AHE: mfg	M.123109552	$\Delta^2 ln$	Avg Hourly Earnings, Prod/Nonsup: Manufacturing

Table B.4: Labor Market

No.	Gp	Short Name	Code	Tran	Descripton
*51	3	Starts: nonfarm	M_110155536	Δln	Housing Starts: Nonfarm(1947-58); Total Farm& Nonfarm(1959-)
*52	3	Starts: NE	M_110155538	Δln	Housing Starts:Northeast
*53	3	Starts: MW	M_110155537	Δln	Housing Starts:Midwest
*54	3	Starts: South	M_110155543	Δln	Housing Starts:South
*55	3	Starts: West	M_110155544	Δln	Housing Starts:West
*56	3	BP: total	M_110155532	Δln	Housing Authorized: Total New Priv Housing Units
*57	3	BP: NE	M_110155531	Δln	Houses Authorized By Build. Permits:Northeast
*58	3	BP: MW	M_110155530	Δln	Houses Authorized By Build. Permits:Midwest
*59	3	BP: South	M_110155533	Δln	Houses Authorized By Build. Permits:South
*60	3	BP: West	M_110155534	Δln	Houses Authorized By Build. Permits:West

Table B.5: Housing

No.	Gp	Short Name	Code	Tran	Descripton
61	4	PMI	M_110157208	lv	Purchasing Managers' Index
62	4	NAPM new ordrs	M_110157210	lv	Napm New Orders Index
63	4	NAPM vendor del	M_110157205	lv	Napm Vendor Deliveries Index
64	4	NAPM Invent	M_110157211	lv	Napm Inventories Index
65	4	Orders: cons gds	M_14385863	Δln	Mfrs' New Orders, Consumer Goods And Materials
66	4	Orders: dble gds	M_14386110	Δln	Mfrs' New Orders, Durable Goods Industries
67	4	Orders: cap gds	M_178554409	Δln	Mfrs' New Orders, Nondefense Capital Goods
68	4	Unf orders: dble	M_14385946	Δln	Mfrs' Unfilled Orders, Durable Goods Indus.
69	4	M&T invent	M_15192014	Δln	Manufacturing And Trade Inventories
70	4	M&T invent/sales	M_15191529	Δlv	Ratio, Mfg. And Trade Inventories To Sales
3	4	Consumption	M_123008274	Δln	Real Personal Consumption Expenditures (AC)
4	4	M&T sales	M_110156998	Δln	Manufacturing And Trade Sales
5	4	Retail sales	M_130439509	Δln	Sales Of Retail Stores
132	4	Consumer expect	hhsntn	Δlv	U. Of Mich. Index Of Consumer Expectations (UM)

Table B.6: Consumption, Orders, and Inventories

No.	Gp	Short Name	Code	Tran	Descripton
71	5	M1	M_110154984	$\Delta^2 ln$	Money Stock: M1
72	5	M2	M_110154985	$\Delta^2 ln$	Money Stock: M2
73	5	Currency	M_110155013	$\Delta^2 ln$	Money Stock: Currency held by the public
74	5	M2 (real)	M_110154985	Δln	Money Supply: Real M2 (AC)
75	5	MB	M_110154995	$\Delta^2 ln$	Monetary Base, Adj For Reserve Requirement Changes
76	5	Reserves tot	M_110155011	$\Delta^2 ln$	Depository Inst Reserves:Total, Adj For Reserve Req Chgs
77	5	Reserves nonbor	M_110155009	$\Delta^2 ln$	Depository Inst Reserves:Nonborrowed,Adj Res Req Chgs
78	5	C&I loans	BUSLOANS	$\Delta^2 ln$	Comm/Indus Loans: Commercial Banks (FRED)
79	5	C&I loans	BUSLOANS	lv	Change in Comm/Indus Loans: Commercial Banks (FRED)
80	5	Cons credit	M_110155009	$\Delta^2 ln$	Consumer Credit Outstanding - Nonrevolving
81	5	Inst cred/PI	M_110154569	Δlv	Ratio, Consumer Installment Credit To Personal Income

Table B.7: Money and Credit

No.	Gp	Short Name	Code	Tran	Description
86	6	Fed Funds	M_110155157	Δlv	Interest Rate: Federal Funds
87	6	Comm paper	CPF3M	Δlv	3-Month AA Financial Commercial Paper Rate (FRED)
88	6	3 mo T-bill	M_110155165	Δlv	Interest Rate: U.S.Treasury Bills,Sec Mkt,3-Mo.
89	6	6 mo T-bill	M_110155165	Δlv	Interest Rate: U.S.Treasury Bills,Sec Mkt,6-Mo.
90	6	1 yr T-bond	M_110155165	Δlv	Interest Rate: U.S.Treasury Const Maturities,1-Yr.
91	6	5 yr T-bond	M_110155174	Δlv	Interest Rate: U.S.Treasury Const Maturities,5-Yr.
92	6	10 yr T-bond	M_110155169	Δlv	Interest Rate: U.S.Treasury Const Maturities,10-Yr.
93	6	Aaa bond	M_14386682	Δlv	Bond Yield: Moody's Aaa Corporate
94	6	Baa bond	M_14386683	Δlv	Bond Yield: Moody's Baa Corporate
95	6	CP-FF spread	-	lv	CP-FF spread (AC)
96	6	3 mo-FF spread	-	lv	3 mo-FF spread (AC)
97	6	6 mo-FF spread	-	lv	6 mo-FF spread (AC)
98	6	1 yr-FF spread	-	lv	1 yr-FF spread (AC)
99	6	5 yr-FF spread	-	lv	5 yr-FF spread (AC)
100	6	10 yr-FF spread	-	lv	10 yr-FF spread (AC)
101	6	Aaa-FF spread	-	lv	Aaa-FF spread (AC)
102	6	Baa-FF spread	-	lv	Baa-FF spread (AC)
103	6	Ex rate: avg	-	Δln	Nominal Effective Exchange Rate, Unit Labor Costs (IMF)
104	6	Ex rate: Switz	M_110154768	Δln	Foreign Exchange Rate: Switzerland - Swiss Franc Per U.S.\$
105	6	Ex rate: Japan	M_110154768	Δln	Foreign Exchange Rate: Japan - Yen Per U.S.\$
106	6	Ex rate: UK	M_110154772	Δln	Foreign Exchange Rate: United Kingdom - Cents Per Pound
107	6	EX rate: Canada	M_110154744	Δln	Foreign Exchange Rate: Canada - Canadian \$ Per U.S.\$

Table B.8: Bond and Exchange Rates

No.	Gp	Short Name	Code	Tran	Description
108	7	PPI: fin gds	M_110157517	$\Delta^2 ln$	Prod. Price Ind.: Finished Goods
109	7	PPI: cons gds	M_110157508	$\Delta^2 ln$	Prod. Price Ind.: Finished Consumer Goods
110	7	PPI: int materials	M_110157527	$\Delta^2 ln$	Prod. Price Ind.: Int. Mat. Supplies & Components
111	7	PPI: crude materials TM ls	M_110157500	$\Delta^2 ln$	Prod. Price Ind.: Crude Materials
112	7	Spot market price	M_110157273	$\Delta^2 ln$	Spot mkt price index: BLS & CRB: all commodities
113	7	PPI: nonferrous materials	M_110157335	$\Delta^2 ln$	Prod. Price Ind.: Nonferrous Materials
114	7	NAPM com price	M_110157204	lv	Napm Commodity Prices Index
115	7	CPI-U: all	M_110157323	$\Delta^2 ln$	Cpi-U: All Items
116	7	CPI-U: apparel	M_110157299	$\Delta^2 ln$	Cpi-U: Apparel & Upkeep
117	7	CPI-U: transp	M_110157302	$\Delta^2 ln$	Cpi-U: Transportation
118	7	CPI-U: medical	M_110157304	$\Delta^2 ln$	Cpi-U: Medical Care
119	7	CPI-U: comm.	M_110157314	$\Delta^2 ln$	Cpi-U: Commodities
120	7	CPI-U: dbles	M_110157315	$\Delta^2 ln$	Cpi-U: Durables
121	7	CPI-U: services	M_110157325	$\Delta^2 ln$	Cpi-U: Services
122	7	CPI-U: ex food	M_110157328	$\Delta^2 ln$	Cpi-U: All Items Less Food
123	7	CPI-U: ex shelter	M_110157329	$\Delta^2 ln$	Cpi-U: All Items Less Shelter
124	7	CPI-U: ex med	M_110157330	$\Delta^2 ln$	Cpi-U: All Items Less Midical Care
125	7	PCE defl	gmdc	$\Delta^2 ln$	Pce, Impl Pr Defl:Pce (BEA)
126	7	PCE defl: dlbes	gmdd	$\Delta^2 ln$	Pce, Impl Pr Defl:Pce; Durables (BEA)
127	7	PCE defl: nondble	gmddn	$\Delta^2 ln$	Pce, Impl Pr Defl:Pce; Nondurables (BEA)
128	7	PCE defl: service	gmddcs	$\Delta^2 ln$	Pce, Impl Pr Defl:Pce; Services (BEA)

Table B.9: Prices

No.	Gp	Short Name	Code	Tran	Description
82	8	S&P 500	M_110155044	Δln	S&P's Common Stock Price Index: Composite
83	8	S&P: indust	M_110155047	Δln	S&P's Common Stock Price Index: & Industrials
84	8	S&P div yield	-	Δlv	S&P's Composite Common Stock: Dividend Yield Real (S)
85	8	S&P PE ratio	-	Δln	S&P's Composite Common Stock: Price-Earnings Ratio Real (S)

Table B.10: Stock Market

B.2.2 Financial Dataset

The data set is at monthly frequency, with 147 observations spanning the period 1960:01-2013:01. All returns and spreads are expressed in logs (i.e. the log of the gross return or spread), are displayed in percent (i.e. multiplied by 100), and are annualized by multiplying by 12, i.e., if x is the original return or spread, we transform to $1200 \ln(1 + x/100)$. Federal Reserve data are annualized by default and are therefore not “re-annualized.” Note: this annualization means that the annualized standard deviation (volatility) is equal to the data standard deviation divided by $\sqrt{12}$. The data series used in this dataset are listed by data source in Table B.11. Additional details on data transformations are given here.

CRSP Data Details

Value-weighted price and dividend data were obtained from the Center for Research in Security Prices (CRSP). From the Annual Update data, we obtain monthly value-weighted returns series *vwretd* (with dividends) and *vwretx* (excluding dividends). These series have the interpretation

$$VWRET D_t = \frac{P_{t+1} + D_{t+1}}{P_t}$$
$$VWRET X_t = \frac{P_{t+1}}{P_t}$$

From these series, a normalized price series P , can be constructed using the recursion

$$P_0 = 1$$
$$P_t = P_{t-1} \cdot VWRET X_t.$$

A dividend series can then be constructed using

$$D_t = P_{t-1}(VWRET D_t - VWRET X_t).$$

No.	Short Name	Source	Tran	Description
1	D.log(DIV)	CRSP	$\Delta \ln$	$\Delta \log D_t^*$ see additional details below
2	D.log(P)	CRSP	$\Delta \ln$	$\Delta \log P_t$ see additional details below
3	D.DIVreinvest	CRSP	$\Delta \ln$	$\Delta \log D_t^{re,*}$ see additional details below
4	D.Preinvest	CRSP	$\Delta \ln$	$\Delta \log P_t^{re,*}$ see additional details below
5	d-p	CRSP	\ln	$\log(D_t^*) - \log P_t$ see additional details below
6	R15-R11	Kenneth French	lv	(Small, High) minus (Small, Low) sorted on (size, btm)
7	CP	Monika Piazzesi	lv	Cochrane-Piazzesi factor (Cochrane, 2005)
8	Mkt-RF	Kenneth French	lv	Market excess return
9	SMB	Kenneth French	lv	Small Minus Big, sorted on size
10	HML	Kenneth French	lv	High Minus Low, sorted on book-to-market
11	UMD	Kenneth French	lv	Up Minus Down, sorted on momentum
12	Agric	Kenneth French	lv	Agric industry portfolio
13	Food	Kenneth French	lv	Food industry portfolio
14	Beer	Kenneth French	lv	Beer industry portfolio
15	Smoke	Kenneth French	lv	Smoke industry portfolio
16	Toys	Kenneth French	lv	Toys industry portfolio
17	Fun	Kenneth French	lv	Fun industry portfolio
18	Books	Kenneth French	lv	Books industry portfolio
19	Hshld	Kenneth French	lv	Hshld industry portfolio
20	Clths	Kenneth French	lv	Clths industry portfolio
21	MedEq	Kenneth French	lv	MedEq industry portfolio
22	Drugs	Kenneth French	lv	Drugs industry portfolio
23	Chems	Kenneth French	lv	Chems industry portfolio
24	Rubbr	Kenneth French	lv	Rubbr industry portfolio
25	Txtls	Kenneth French	lv	Txtls industry portfolio
26	BldMt	Kenneth French	lv	BldMt industry portfolio
27	Cnstr	Kenneth French	lv	Cnstr industry portfolio
28	Steel	Kenneth French	lv	Steel industry portfolio
39	Mach	Kenneth French	lv	Mach industry portfolio
30	ElcEq	Kenneth French	lv	ElcEq industry portfolio
31	Autos	Kenneth French	lv	Autos industry portfolio
32	Aero	Kenneth French	lv	Aero industry portfolio
33	Ships	Kenneth French	lv	Ships industry portfolio
34	Mines	Kenneth French	lv	Mines industry portfolio
35	Coal	Kenneth French	lv	Coal industry portfolio
36	Oil	Kenneth French	lv	Oil industry portfolio
37	Util	Kenneth French	lv	Util industry portfolio
38	Telcm	Kenneth French	lv	Telcm industry portfolio
39	PerSv	Kenneth French	lv	PerSv industry portfolio
40	BusSv	Kenneth French	lv	BusSv industry portfolio
41	Hardw	Kenneth French	lv	Hardw industry portfolio
42	Chips	Kenneth French	lv	Chips industry portfolio
43	LabEq	Kenneth French	lv	LabEq industry portfolio
44	Paper	Kenneth French	lv	Paper industry portfolio
45	Boxes	Kenneth French	lv	Boxes industry portfolio
46	Trans	Kenneth French	lv	Trans industry portfolio
47	Whlsl	Kenneth French	lv	Whlsl industry portfolio
48	Rtail	Kenneth French	lv	Rtail industry portfolio
49	Meals	Kenneth French	lv	Meals industry portfolio
50	Banks	Kenneth French	lv	Banks industry portfolio
51	Insur	Kenneth French	lv	Insur industry portfolio
52	RIEst	Kenneth French	lv	RIEst industry portfolio
53	Fin	Kenneth French	lv	Fin industry portfolio
54	Other	Kenneth French	lv	Other industry portfolio

Table B.11: List of variables in the financial dataset.

List of variables in the financial dataset (continued)

No.	Short Name	Source	Tran	Description
55	1.2	Kenneth French	<i>lv</i>	(1, 2) portfolio sorted on (size, book-to-market)
56	1.4	Kenneth French	<i>lv</i>	(1, 4) portfolio sorted on (size, book-to-market)
57	1.5	Kenneth French	<i>lv</i>	(1, 5) portfolio sorted on (size, book-to-market)
58	1.6	Kenneth French	<i>lv</i>	(1, 6) portfolio sorted on (size, book-to-market)
59	1.7	Kenneth French	<i>lv</i>	(1, 7) portfolio sorted on (size, book-to-market)
60	1.8	Kenneth French	<i>lv</i>	(1, 8) portfolio sorted on (size, book-to-market)
61	1.9	Kenneth French	<i>lv</i>	(1, 9) portfolio sorted on (size, book-to-market)
62	1.high	Kenneth French	<i>lv</i>	(1, high) portfolio sorted on (size, book-to-market)
63	2.low	Kenneth French	<i>lv</i>	(2, low) portfolio sorted on (size, book-to-market)
64	2.2	Kenneth French	<i>lv</i>	(2, 2) portfolio sorted on (size, book-to-market)
65	2.3	Kenneth French	<i>lv</i>	(2, 3) portfolio sorted on (size, book-to-market)
66	2.4	Kenneth French	<i>lv</i>	(2, 4) portfolio sorted on (size, book-to-market)
67	2.5	Kenneth French	<i>lv</i>	(2, 5) portfolio sorted on (size, book-to-market)
68	2.6	Kenneth French	<i>lv</i>	(2, 6) portfolio sorted on (size, book-to-market)
69	2.7	Kenneth French	<i>lv</i>	(2, 7) portfolio sorted on (size, book-to-market)
70	2.8	Kenneth French	<i>lv</i>	(2, 8) portfolio sorted on (size, book-to-market)
71	2.9	Kenneth French	<i>lv</i>	(2, 9) portfolio sorted on (size, book-to-market)
72	2.high	Kenneth French	<i>lv</i>	(2, high) portfolio sorted on (size, book-to-market)
73	3.low	Kenneth French	<i>lv</i>	(3, low) portfolio sorted on (size, book-to-market)
74	3.2	Kenneth French	<i>lv</i>	(3, 2) portfolio sorted on (size, book-to-market)
75	3.3	Kenneth French	<i>lv</i>	(3, 3) portfolio sorted on (size, book-to-market)
76	3.4	Kenneth French	<i>lv</i>	(3, 4) portfolio sorted on (size, book-to-market)
77	3.5	Kenneth French	<i>lv</i>	(3, 5) portfolio sorted on (size, book-to-market)
78	3.6	Kenneth French	<i>lv</i>	(3, 6) portfolio sorted on (size, book-to-market)
79	3.7	Kenneth French	<i>lv</i>	(3, 7) portfolio sorted on (size, book-to-market)
80	3.8	Kenneth French	<i>lv</i>	(3, 8) portfolio sorted on (size, book-to-market)
81	3.9	Kenneth French	<i>lv</i>	(3, 9) portfolio sorted on (size, book-to-market)
82	3.high	Kenneth French	<i>lv</i>	(3, high) portfolio sorted on (size, book-to-market)
83	4.low	Kenneth French	<i>lv</i>	(4, low) portfolio sorted on (size, book-to-market)
84	4.2	Kenneth French	<i>lv</i>	(4, 2) portfolio sorted on (size, book-to-market)
85	4.3	Kenneth French	<i>lv</i>	(4, 3) portfolio sorted on (size, book-to-market)
86	4.4	Kenneth French	<i>lv</i>	(4, 4) portfolio sorted on (size, book-to-market)
87	4.5	Kenneth French	<i>lv</i>	(4, 5) portfolio sorted on (size, book-to-market)
88	4.6	Kenneth French	<i>lv</i>	(4, 6) portfolio sorted on (size, book-to-market)
89	4.7	Kenneth French	<i>lv</i>	(4, 7) portfolio sorted on (size, book-to-market)
90	4.8	Kenneth French	<i>lv</i>	(4, 8) portfolio sorted on (size, book-to-market)
91	4.9	Kenneth French	<i>lv</i>	(4, 9) portfolio sorted on (size, book-to-market)
92	4.high	Kenneth French	<i>lv</i>	(4, high) portfolio sorted on (size, book-to-market)
93	5.low	Kenneth French	<i>lv</i>	(5, low) portfolio sorted on (size, book-to-market)
94	5.2	Kenneth French	<i>lv</i>	(5, 2) portfolio sorted on (size, book-to-market)
95	5.3	Kenneth French	<i>lv</i>	(5, 3) portfolio sorted on (size, book-to-market)
96	5.4	Kenneth French	<i>lv</i>	(5, 4) portfolio sorted on (size, book-to-market)
97	5.5	Kenneth French	<i>lv</i>	(5, 5) portfolio sorted on (size, book-to-market)
98	5.6	Kenneth French	<i>lv</i>	(5, 6) portfolio sorted on (size, book-to-market)
99	5.7	Kenneth French	<i>lv</i>	(5, 7) portfolio sorted on (size, book-to-market)
100	5.8	Kenneth French	<i>lv</i>	(5, 8) portfolio sorted on (size, book-to-market)
101	5.9	Kenneth French	<i>lv</i>	(5, 9) portfolio sorted on (size, book-to-market)
102	5.high	Kenneth French	<i>lv</i>	(5, high) portfolio sorted on (size, book-to-market)

List of variables in the financial dataset (continued)

No.	Short Name	Source	Tran	Description
103	6_low	Kenneth French	<i>lv</i>	(6, low) portfolio sorted on (size, book-to-market)
104	6_2	Kenneth French	<i>lv</i>	(6, 2) portfolio sorted on (size, book-to-market)
105	6_3	Kenneth French	<i>lv</i>	(6, 3) portfolio sorted on (size, book-to-market)
106	6_4	Kenneth French	<i>lv</i>	(6, 4) portfolio sorted on (size, book-to-market)
107	6_5	Kenneth French	<i>lv</i>	(6, 5) portfolio sorted on (size, book-to-market)
108	6_6	Kenneth French	<i>lv</i>	(6, 6) portfolio sorted on (size, book-to-market)
109	6_7	Kenneth French	<i>lv</i>	(6, 7) portfolio sorted on (size, book-to-market)
110	6_8	Kenneth French	<i>lv</i>	(6, 8) portfolio sorted on (size, book-to-market)
111	6_9	Kenneth French	<i>lv</i>	(6, 9) portfolio sorted on (size, book-to-market)
112	6_high	Kenneth French	<i>lv</i>	(6, high) portfolio sorted on (size, book-to-market)
113	7_low	Kenneth French	<i>lv</i>	(7, low) portfolio sorted on (size, book-to-market)
114	7_2	Kenneth French	<i>lv</i>	(7, 2) portfolio sorted on (size, book-to-market)
115	7_3	Kenneth French	<i>lv</i>	(7, 3) portfolio sorted on (size, book-to-market)
116	7_4	Kenneth French	<i>lv</i>	(7, 4) portfolio sorted on (size, book-to-market)
117	7_5	Kenneth French	<i>lv</i>	(7, 5) portfolio sorted on (size, book-to-market)
118	7_6	Kenneth French	<i>lv</i>	(7, 6) portfolio sorted on (size, book-to-market)
119	7_7	Kenneth French	<i>lv</i>	(7, 7) portfolio sorted on (size, book-to-market)
120	7_8	Kenneth French	<i>lv</i>	(7, 8) portfolio sorted on (size, book-to-market)
121	7_9	Kenneth French	<i>lv</i>	(7, 9) portfolio sorted on (size, book-to-market)
122	8_low	Kenneth French	<i>lv</i>	(8, low) portfolio sorted on (size, book-to-market)
123	8_2	Kenneth French	<i>lv</i>	(8, 2) portfolio sorted on (size, book-to-market)
124	8_3	Kenneth French	<i>lv</i>	(8, 3) portfolio sorted on (size, book-to-market)
125	8_4	Kenneth French	<i>lv</i>	(8, 4) portfolio sorted on (size, book-to-market)
126	8_5	Kenneth French	<i>lv</i>	(8, 5) portfolio sorted on (size, book-to-market)
127	8_6	Kenneth French	<i>lv</i>	(8, 6) portfolio sorted on (size, book-to-market)
128	8_7	Kenneth French	<i>lv</i>	(8, 7) portfolio sorted on (size, book-to-market)
129	8_8	Kenneth French	<i>lv</i>	(8, 8) portfolio sorted on (size, book-to-market)
130	8_9	Kenneth French	<i>lv</i>	(8, 9) portfolio sorted on (size, book-to-market)
131	8_high	Kenneth French	<i>lv</i>	(8, high) portfolio sorted on (size, book-to-market)
132	9_low	Kenneth French	<i>lv</i>	(9, low) portfolio sorted on (size, book-to-market)
133	9_2	Kenneth French	<i>lv</i>	(9, 2) portfolio sorted on (size, book-to-market)
134	9_3	Kenneth French	<i>lv</i>	(9, 3) portfolio sorted on (size, book-to-market)
135	9_4	Kenneth French	<i>lv</i>	(9, 4) portfolio sorted on (size, book-to-market)
136	9_5	Kenneth French	<i>lv</i>	(9, 5) portfolio sorted on (size, book-to-market)
137	9_6	Kenneth French	<i>lv</i>	(9, 6) portfolio sorted on (size, book-to-market)
138	9_7	Kenneth French	<i>lv</i>	(9, 7) portfolio sorted on (size, book-to-market)
139	9_8	Kenneth French	<i>lv</i>	(9, 8) portfolio sorted on (size, book-to-market)
140	9_high	Kenneth French	<i>lv</i>	(9, high) portfolio sorted on (size, book-to-market)
141	10_low	Kenneth French	<i>lv</i>	(10, low) portfolio sorted on (size, book-to-market)
142	10_2	Kenneth French	<i>lv</i>	(10, 2) portfolio sorted on (size, book-to-market)
143	10_3	Kenneth French	<i>lv</i>	(10, 3) portfolio sorted on (size, book-to-market)
144	10_4	Kenneth French	<i>lv</i>	(10, 4) portfolio sorted on (size, book-to-market)
145	10_5	Kenneth French	<i>lv</i>	(10, 5) portfolio sorted on (size, book-to-market)
146	10_6	Kenneth French	<i>lv</i>	(10, 6) portfolio sorted on (size, book-to-market)
147	10_7	Kenneth French	<i>lv</i>	(10, 7) portfolio sorted on (size, book-to-market)

We define the series

$$D_t^* = (D_t + D_{t-1} + D_{t-2} + D_{t-3}).$$

For the price and dividend series under “reinvestment,” we calculate the price under reinvestment, P_t^{re} , as the normalized value of the market portfolio under reinvestment of dividends, using the recursion

$$P_0^{re} = 1$$

$$P_t^{re} = P_{t-1} \cdot VWRET D_t$$

Similarly, we can define dividends under reinvestment, D_t^{re} , as the total dividend payments on this portfolio (the number of “shares” of which have increased over time) using

$$D_t^{re} = P_{t-1}^{re} (VWRET D_t - VWRET X_t).$$

As before, we define the series

$$D_t^{re,*} = (D_t^{re} + D_{t-1}^{re} + D_{t-2}^{re} + D_{t-3}^{re})$$

Five data series are constructed from the CRSP data as follows:

- (1) D_log(DIV): $\Delta \log D_t^*$.
- (2) D_log(P): $\Delta \log P_t$.
- (3) D_DIVreinvest: $\Delta \log D_t^{re,*}$
- (4) D_Preinvest: $\Delta \log P_t^{re,*}$
- (5) d-p: $\log(D_t^*) - \log(P_t)$

Kenneth French Data Details

The following data are obtained from the data library of Kenneth French (http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html):

- Fama/French Factors: From this dataset we obtain the data series RF, Mkt-RF, SMB, HML.
- 25 Portfolios formed on Size and Book-to-Market (5 x 5): From this dataset we obtain the series R15-R11, which is the spread between the (small, high book-to-market) and (small, low book-to-market) portfolios.
- Momentum Factor (Mom): From this dataset we obtain the series UMD, which is equal to the momentum factor.
- 49 Industry Portfolios: From this dataset we use all value-weighted series, excluding any series that have missing observations from Jan. 1960 on, from which we obtain the series Agric through Other. The omitted series are: Soda, Hlth, FabPr, Guns, Gold, Softw.
- 100 Portfolios formed in Size and Book-to-Market: From this dataset we use all value-weighted series, excluding any series that have missing observations from Jan. 1960 on. This yields variables with the name X_Y where X stands for the index of the size variable (1, 2, ..., 10) and Y stands for the index of the book-to-market variable (Low, 2, 3, ..., 8, 9, High). The omitted series are 1_low, 1_3, 7_high, 9_9, 10_8, 10_9, 10_high.

B.2.3 Firm-level Dataset

Firm level observations are from COMPUSTAT Fundamentals Quarterly dataset. The unit of observation is the change in firm pre-tax profits $P_{i,t}$, normalized by a two-period moving average of sales, $S_{i,t}$, following Bloom (2009). Bloom constructs

$$dpretax_{i,t} = (P_{it} - P_{it-1}) / (0.5 \cdot S_{it} + 0.5 \cdot S_{it-1}), \quad (\text{B.2})$$

for each firm i in quarter t . This is the same measure reported on in Bloom (2009), Table 1, and discussed in footnote c. We find, however, that (B.2) exhibits clear seasonality patterns, thus we instead use year-over-year changes for the variable (B.2), normalized by average

sales:

$$Y_{i,t} = dpretaxy_{i,t} = (P_{it} - P_{it-4}) / (0.5 \cdot S_{it} + 0.5 \cdot S_{t-4}), \quad (\text{B.3})$$

We follow the trimming procedures used by Bloom, which includes considering any observation with sales $S = 0$ a missing value, and windsorizing observations at the top and bottom 0.05% values (replacing values in the top and bottom 0.05% with the values at the 0.05th and 99.95th percentile values).⁵ After converting to a balanced panel, we are left with 155 firms from 1970:Q1-2011:Q2 without missing values.

These variables are constructed from COMPUSTAT Fundamentals Quarterly dataset. It contains 155 firms observed from 1970Q1 to 2011Q2 that have non-missing observations for $P_{i,t}$ (Compustat identifier *piq*) and $S_{i,t}$ (Compustat identifier for net sales *saleq*) across the entire time period.⁶

- gvkey: firm identifier
- date: period (1 to 166)
- *dpretax*: quarterly change in pretax profits scaled by average sales in current and past quarter:

$$dpretax_{i,t} = \frac{piq_{i,t} - piq_{i,t-1}}{0.5 (saleq_{i,t} + saleq_{i,t-1})}.$$

- *dpretaxy*: year-over-year change in quarterly pretax profits scaled by average sales:

$$dpretaxy_{i,t} = \frac{piq_{i,t} - piq_{i,t-4}}{0.5 (saleq_{i,t} + saleq_{i,t-4})}$$

B.2.4 Data for VAR Analysis

Monthly Macro VAR Endogenous variables, in order:

⁵A detailed description of these procedures are given in the code to Bloom (2009) <http://www.stanford.edu/~nbloom/replication.zip>.

⁶This item represents operating and nonoperating income before provisions for income taxes and minority interest. Earnings (COMPUSTAT code *ibq*) are measured as the income of a company after all expenses, including special items, income taxes, and minority interest, but before provisions for common and/or preferred dividends. Formally: $ibq = piq - txt$ (income taxes) $- mii$ (minority interest).

- (1) $\log(\text{IP})$
- (2) $\log(\text{Employment})$
- (3) $\log(\text{Real Consumption})$
- (4) $\log(\text{Price Level})$
- (5) $\log(\text{Real Value of New Orders})$
- (6) $\log(\text{Real Wage})$
- (7) $\log(\text{Hours})$
- (8) Federal Funds Rate
- (9) $\log(\text{S\&P 500})$
- (10) growth rate of M2
- (11) uncertainty (various measures)

Where

- IP = Industrial Production Index: total; jlnldata series 6.
- Employment = All employees, total nofarm; FRED series PAYEMS.
- Real Consumption = jlnldata series 3.
- Price Level = PCE Implicit Price Deflator; jlnldata series 125.
- New Orders = Value of Manufacturers New Order: consumer goods and materials + Value of Manufacturers' New Orders: nondefense capital goods; jlnldata series 65 + 67.
- Real Value of New Orders = New Orders/Price Level.
- Wage = Average Hourly Earnings of Production and Nonsupervisory Workers: Manufacturing; jlnldata series 131.
- Real Wage = Wage/Price Level.
- Hours = Average Weekly Hours of Production and Nonsupervisory Workers: manufacturing; jlnldata series 49.
- Federal Funds Rate = Effective Federal Funds Rate; jlnldata series 86.

- S&P 500 = jlndata series 82.
- M2 = jlndata series 72.

Monthly Bloom (2009) VAR Endogenous variables, in order:

- (1) $\log(\text{S\&P 500})$
- (2) uncertainty (various measures)
- (3) Federal Funds Rate
- (4) $\log(\text{Nominal Wage})$
- (5) $\log(\text{Price Level})$
- (6) Hours
- (7) $\log(\text{Employment})$
- (8) $\log(\text{Industrial Production})$

Where

- S&P 500 = jlndata series 82.
- Federal Funds Rate = effective federal funds rate; jlndata series 86.
- Nominal Wage = average hourly earnings in manufacturing; jlndata series 131.
- Price Level = CPI-U: all items; jlndata series 115.
- Hours = Average Weekly Hours of Production and Nonsupervisory Workers: manufacturing; jlndata series 49.
- Employment = Employees on Nonfarm Payrolls: manufacturing; jlndata series 37.
- Industrial Production = Industrial Production Index: manufacturing; jlndata series 16.

Appendix C

Appendix for Chapter 3

C.1 Proof of Proposition 1

Defining the variable $\psi_t \equiv \delta_t/D_t$ to denote the fraction of the aggregate dividend held by experts, and substituting out \underline{c}_t from the objective using the resource constraint, the planner's problem can be equivalently rewritten as:

$$\max_{c_t \geq 0, \psi_t \in [0,1]} E \left[\int_0^\infty e^{-\rho t} \{ \tau u(c_t) + (1 - \tau) u((\psi_t a + (1 - \psi_t) \underline{a} + 2\theta) D_t - c_t) \} \right] dt,$$

$$s.t. \quad dD_t = (\psi_t g + (1 - \psi_t) \underline{g}) D_t dt + \sigma D_t dZ_t.$$

The Hamilton-Jacobi-Bellman (HJB) equation associated with this program is

$$\begin{aligned} \rho V(D_t) = & \max_{c_t \geq 0, \psi_t} \{ \tau u(c_t) + (1 - \tau) u((\psi_t a + (1 - \psi_t) \underline{a} + 2\theta) D_t - c_t) \\ & + (\psi_t g + (1 - \psi_t) \underline{g}) D_t V'(D_t) + \frac{1}{2} (\sigma D_t)^2 V''(D_t) \\ & + \vartheta_t^0 D_t \psi_t + \vartheta_t^1 D_t (1 - \psi_t) \}, \end{aligned}$$

where $\vartheta_t^0 D_t, \vartheta_t^1 D_t \geq 0$ are the multipliers associated with the constraints $\psi_t \geq 0$ and $\psi_t \leq 1$, respectively. The first-order necessary conditions for optimality, and the complementary slackness conditions, are:

1. $\tau u'(c_t) = (1 - \tau)u'(\psi_t a + (1 - \psi_t)\underline{a} + 2\theta)D_t - c_t$
2. $(1 - \tau)(a - \underline{a})u'(\psi_t a + (1 - \psi_t)\underline{a} + 2\theta)D_t - c_t + (g - \underline{g})V'(D_t) = \vartheta_t^1 - \vartheta_t^0$
3. $\vartheta_t^0 \psi_t = \vartheta_t^1 (1 - \psi_t) = 0, \vartheta_t^0, \vartheta_t^1 \geq 0$.

Combining the first two optimality conditions,

$$c_t = u'^{-1} \left(\frac{1}{\tau(a - \underline{a})} (\vartheta_t^1 - \vartheta_t^0 - (g - \underline{g})V'(D_t)) \right).$$

Substituting back into the first condition, and solving for ψ_t gives

$$\begin{aligned} \psi_t = & \frac{1}{(a - \underline{a})D_t} \left[u'^{-1} \left(\frac{1}{(1 - \tau)(a - \underline{a})} (\vartheta_t^1 - \vartheta_t^0 - (g - \underline{g})V'(D_t)) \right) \right. \\ & \left. + u'^{-1} \left(\frac{1}{\tau(a - \underline{a})} (\vartheta_t^1 - \vartheta_t^0 - (g - \underline{g})V'(D_t)) \right) - (\underline{a} + 2\theta)D_t \right]. \end{aligned}$$

Next, recall that standard dynamic programming results imply that the value function is concave and strictly increasing because of the form of the objective and the assumptions on $u(\cdot)$; therefore $V'(D_t) > 0$. If neither constraint on ψ_t is binding, then $\vartheta_t^0 = \vartheta_t^1 = 0$ and $\psi_t < 0$ by $u'^{-1}(x) < 0$ for $x < 0$, which is a contradiction. If only the second constraint does not bind, $\vartheta_t^1 = 0$ and $\vartheta_t^0 \geq 0$, it still follows that $\psi_t < 0$. Therefore, the case that only the second constraint binds, $\vartheta_t^0 = 0$, must be optimal. This implies that $\psi_t = 1$. Substituting back into the policy rule for consumption implies that $c_t = \tau(a + 2\theta)D_t$. Using the resource constraint and the definition of ψ_t gives the final expressions presented in the proposition.

C.2 Proof of Proposition 2

The proof of this proposition follows from the assumption of logarithmic preferences, and the assumption that agents do not internalize the effects of their exogenous income stream on net worth. To demonstrate how this is the case, I closely follow Merton (1969, 1971). Let the return on a risky and risk-free security, respectively, be given by

$$dR_t = \mu_t dt + \sigma_t dZ_t, \quad dr_t = r_t dt.$$

Consider the program

$$\max_{c_t, \pi_t} E \int_0^\infty e^{-\rho t} \log(c_t) dt \quad s.t.$$

$$dn_t = (\pi_t(\mu_t - r_t)n_t + r_t n_t - c_t) dt + \pi_t \sigma_t n_t dZ_t$$

$$n_t \geq 0.$$

The Hamilton-Jacobi-Bellman (HJB) equation associated with this program is:

$$\begin{aligned} \rho V(n_t) = & \max_{c_t, \pi_t} \left\{ \log(c_t) + [\pi_t(\mu_t - r_t)n_t + r_t n_t - c_t] V'(n_t) \right. \\ & \left. + \frac{1}{2} (\pi_t \sigma_t n_t)^2 V''(n_t) \right\}. \end{aligned}$$

The first-order necessary conditions for optimality are:

1. $c_t = 1/V'(n_t)$,
2. $\pi_t = - \left(\frac{\mu_t - r_t}{\sigma_t^2} \right) \frac{V'(n_t)}{V''(n_t)n_t}.$

Substituting these back into the HJB equation gives the concentrated (deterministic) second-order differential equation

$$\rho V(n_t) = -\log(V(n_t)) - \frac{1}{2} \left(\frac{\mu_t - r_t}{\sigma_t} \right)^2 \frac{V'(n_t)^2}{V''(n_t)} + r_t n_t V'(n_t) - 1.$$

Guess that the value function has the form $V(n_t) = \frac{1}{\rho}(\log(n_t) + \phi_t)$, so that $V'(n_t) = 1/(\rho n_t)$ and $V''(n_t) = -1/(\rho n_t^2)$. Substituting these into the previous differential equation implies that

$$\phi_t = \log(\rho) + \frac{1}{2\rho} \left(\frac{\mu_t - r_t}{\sigma_t} \right)^2 + \frac{r_t}{\rho} - 1,$$

which indeed does not depend on the state variable n_t . Therefore, the value function is

$$V(n_t) = \frac{1}{\rho} \left[\log(n_t) + \log(\rho) + \frac{1}{2\rho} \left(\frac{\mu_t - r_t}{\sigma_t} \right)^2 + \frac{r_t}{\rho} - 1 \right],$$

and the optimal policy rules are

$$\begin{aligned} c_t &= \rho n_t \\ \pi_t &= \left(\frac{\mu_t - r_t}{\sigma_t^2} \right). \end{aligned}$$

Moreover, if the restriction $\pi_t \geq 0$ is imposed, then $\pi_t = 0$ whenever $\mu_t < r_t$. Letting $n_t = N_t$, $\pi_t = q_t \delta_t / N_t$, $\mu_t - r_t = E_t(dr_t^D - dr_t)$, and $\sigma_t = (\sigma + \sigma_t^q)$ gives the desired result.

C.3 Proof of Proposition 3

The equilibrium of this model features two regimes, due to the no short sale constraint $\delta_t, \underline{\delta}_t \geq 0$. I therefore consider first the regime where it is not optimal for households to finance any of the risky projects ($\psi_t = 1$). Market clearing then implies that

$$q_t = \bar{q} = \frac{a + 2\theta}{\rho}$$

so $\mu_t^q = \sigma_t^q = 0$. In order for households to refrain from lending, it must be that they earn a large enough return on funds deposited with the intermediary:

$$r_t > \frac{\underline{a}\rho}{a + 2\theta} + \underline{g}.$$

In this regime, the expert portfolio optimality condition is, from Proposition 2,

$$\frac{\psi_t}{\eta_t} = \frac{a/q_t + g + \mu_t^q + \sigma\sigma_t^q - r_t}{(\sigma + \sigma_t)^2}.$$

Since prices are constant, this condition requires that

$$r_t = \rho + g - \frac{\sigma^2}{\eta_t}.$$

Lastly, substituting in $\psi_t = 1$, $\mu_t^q = \sigma_t^q = 0$ into the law of motion for η_t (derived in Section 3.2), I find that

$$d\eta_t = \left[\frac{(1 - \eta_t)^2}{\eta_t} \sigma^2 + (a/\bar{q} - \rho)\eta_t + \theta/\bar{q} \right] dt + (1 - \eta_t)\sigma dZ_t.$$

Plugging in the value of \bar{q} delivers the results of the second half of the proposition.

In the other regime, where households find it optimal to directly finance some of the risky projects ($\psi_t < 1$), the evolution of η_t takes the more general form derived in Section 3.2, which I reproduce here for convenience:

$$\begin{aligned} d\eta_t = & \left[(\psi_t - \eta_t)(\mu_t^q + \sigma\sigma_t^q - (\sigma + \sigma_t^q)^2) + \psi_t(a/q_t + g) + \theta/q_t \right. \\ & \left. - (\rho + g\psi_t + \underline{g}(1 - \psi_t))\eta_t \right] dt - (\psi_t - \eta_t)dr_t + (\psi_t - \eta_t)(\sigma + \sigma_t^q)dZ_t. \end{aligned}$$

Since households are invested in the projects, their optimality condition must hold with equality, which from Proposition 2 means that

$$\frac{1 - \psi_t}{1 - \eta_t} = \frac{\underline{a}/q_t + \underline{g} + \mu_t^q + \sigma\sigma_t^q - r_t}{(\sigma + \sigma_t)^2}.$$

Solving this expression for $\mu_t^q + \sigma\sigma_t^q - r_t$ and substituting into the drift term of $d\eta_t$ above gives the desired law of motion for the state variable. Furthermore, since expert optimality

also must hold whenever $\psi_t > 0$, I can equate household and expert portfolio rules to obtain

$$(a - \underline{a})/q_t + (g - \underline{g}) = (\sigma + \sigma_t^2) \left(\frac{\psi_t}{\eta_t} - \frac{1 - \psi_t}{1 - \eta_t} \right).$$

Solving this expression for σ_t^q gives the solution for endogenous price volatility. The only remaining step is to characterize the equilibrium behavior of the price process as a function of the state variable, $q_t = q(\eta_t)$. By Itô's Lemma,

$$\begin{aligned} dq(\eta_t) &= d\eta_t q'(\eta_t) + \frac{1}{2}(\sigma_t^\eta)^2 q''(\eta_t) dt \\ &= \underbrace{[\mu_t^\eta q'(\eta_t) + \frac{1}{2}(\sigma_t^\eta)^2 q''(\eta_t)] dt}_{\mu_t^q q_t} + \underbrace{[\sigma_t^\eta q'(\eta_t)] dZ_t}_{\sigma_t^q q_t} \end{aligned}$$

Matching coefficients with the law of motion $dq_t/q_t = \mu_t^q dt + \sigma_t^q dZ_t$, and using the law of motion for η_t above gives the remaining functions. After finding all of the other price coefficients, I retrieve r_t in this regime from the households optimality condition above.

C.4 Proof of Proposition 4

Assume that the law of motion of individual expert net worth follows

$$dN_t = \mu_t^N dt + \sigma_t^N dZ_t.$$

Index time as $t, t + \Delta t, t + 2\Delta t, \dots$ and define $\Delta N_t \equiv (N_{t+\Delta t} - N_t)/\sqrt{\Delta t}$ to be the rescaled change in the value of net worth from t to $t + \Delta t$. By the Brownian nature of the driving process $\{Z_t\}$,

$$\Delta N_t \sim \mathcal{N} \left(\mu_t^N N_t \sqrt{\Delta t}, (\sigma_t^N N_t)^2 \right).$$

Definition 3. *Given a tolerance probability $\pi_v \in (0, 1)$, the agent's rescaled VaR from period*

t to $t + \Delta t$ is:

$$v_t \equiv \inf\{v \geq 0 \mid P(\Delta N_t \leq -v) = \pi_v\}.$$

This says that v_t is the minimum amount of equity needed to ensure that losses greater than v_t per $\sqrt{\Delta t}$ units of time only happen with a small probability π_v . A general VaR constraint requires that the agent hold a multiple $\beta \geq 0$ of his net worth to cover his Value-at-Risk:

$$v_t \leq \beta N_t.$$

By the normality assumption, $P(\Delta N_t \leq -v) = \pi_v$ holds if and only if

$$\begin{aligned} \Phi\left(\frac{-v - \mu_t^N \sqrt{\Delta t}}{\sigma_t^N}\right) &= \pi_v, \Leftrightarrow \\ \frac{v + \mu_t^N \sqrt{\Delta t}}{\sigma_t^N} &= \Phi^{-1}(1 - \pi_v), \Leftrightarrow \\ v &= \Phi^{-1}(1 - \pi_v)\sigma_t^N - \mu_t^N \sqrt{\Delta t}. \end{aligned}$$

Now, letting $\Delta t \rightarrow 0$, it follows that the agents Value-at-Risk takes the form

$$v_t = \Phi^{-1}(1 - \pi_v)\sigma_t^N.$$

Combining with the constraint that $v_t \leq \beta N_t$, defining $\alpha \equiv \Phi^{-1}(1 - \pi_v)/\beta$, and noting that $\sigma_t^N = q_t \delta_t (\sigma + \sigma_t^q)$ gives the desired result.

C.5 Proof of Proposition 5

I proceed in the same way as in Proposition 3. Let the return on a risky and risk-free security, respectively, be given by

$$dR_t = \mu_t dt + \sigma_t dZ_t, \quad dr_t = r_t dt.$$

Consider the program

$$\begin{aligned} \max_{c_t, \pi_t} E \int_0^\infty e^{-\rho t} \log(c_t) dt \quad s.t. \\ dn_t = (\pi_t(\mu_t - r_t)n_t + r_t n_t - c_t) dt + \pi_t \sigma_t n_t dZ_t \\ \alpha \pi_t \sigma_t \leq 1 \\ n_t \geq 0. \end{aligned}$$

The Hamilton-Jacobi-Bellman (HJB) equation associated with this program is:

$$\begin{aligned} \rho V(n_t) = \max_{c_t, \pi_t} \{ \log(c_t) + [\pi_t(\mu_t - r_t)n_t + r_t n_t - c_t] V'(n_t) \\ + \frac{1}{2} (\pi_t \sigma_t n_t)^2 V''(n_t) + \varphi_t (1 - \alpha \pi_t \sigma_t) \}, \end{aligned}$$

where $\varphi_t \geq 0$ is the multiplier on the volatility constraint. The first-order necessary conditions for optimality, and the complementary slackness conditions, are:

1. $c_t = 1/V'(n_t)$,
2. $\pi_t = - \left(\frac{\mu_t - r_t}{\sigma_t^2} \right) \frac{V'(n_t)}{V''(n_t)n_t} + \alpha \frac{\varphi_t}{\sigma_t n_t^2 V''(n_t)}$,
3. $\varphi_t (1 - \alpha \pi_t \sigma_t) = 0$.

The second and third optimality conditions can be written as

$$\pi_t = \min \left\{ - \left(\frac{\mu_t - r_t}{\sigma_t^2} \right) \frac{V'(n_t)}{V''(n_t)n_t}, \frac{1}{\alpha \sigma_t} \right\}.$$

Substituting these back into the HJB equation gives the concentrated (deterministic) second-order differential equation

$$\begin{aligned}\rho V(n_t) &= -\log(V(n_t)) + \min \left\{ -\left(\frac{\mu_t - r_t}{\sigma_t^2} \right) \frac{V'(n_t)}{V''(n_t)n_t}, \frac{1}{\alpha\sigma_t} \right\} \\ &\quad \times (\mu_t - r_t)n_t V'(n_t) + r_t n_t V'(n_t) - 1 \\ &\quad + \frac{1}{2} \min \left\{ -\left(\frac{\mu_t - r_t}{\sigma_t^2} \right) \frac{V'(n_t)}{V''(n_t)n_t}, \frac{1}{\alpha\sigma_t} \right\}^2 (\sigma_t n_t)^2 V''(n_t).\end{aligned}$$

Guess that the value function has the form $V(n_t) = \frac{1}{\rho}(\log(n_t) + \phi_t)$, so that $V'(n_t) = 1/(\rho n_t)$ and $V''(n_t) = -1/(\rho n_t^2)$. Substituting these into the previous differential equation implies that

$$\begin{aligned}\phi_t &= \log(\rho) + \min \left\{ \frac{\mu_t - r_t}{\sigma_t^2}, \frac{1}{\alpha\sigma_t} \right\} \frac{(\mu_t - r_t)}{\rho} + \frac{r_t}{\rho} - 1 \\ &\quad - \frac{1}{2\rho} \min \left\{ \frac{\mu_t - r_t}{\sigma_t^2}, \frac{1}{\alpha\sigma_t} \right\}^2 \sigma_t^2.\end{aligned}$$

which indeed does not depend on the state variable n_t . Therefore, the value function is

$$\begin{aligned}V(n_t) &= \frac{1}{\rho} \left[\log(n_t) + \log(\rho) + \min \left\{ \frac{\mu_t - r_t}{\sigma_t^2}, \frac{1}{\alpha\sigma_t} \right\} \frac{(\mu_t - r_t)}{\rho} \right. \\ &\quad \left. + \frac{r_t}{\rho} - 1 - \frac{1}{2\rho} \min \left\{ \frac{\mu_t - r_t}{\sigma_t^2}, \frac{1}{\alpha\sigma_t} \right\}^2 \sigma_t^2 \right]\end{aligned}$$

and the optimal policy rules are

$$\begin{aligned}c_t &= \rho n_t, \\ \pi_t &= \min \left\{ \frac{\mu_t - r_t}{\sigma_t^2}, \frac{1}{\alpha\sigma_t} \right\}.\end{aligned}$$

Moreover, if the restriction $\pi_t \geq 0$ is imposed, then $\pi_t = 0$ whenever $\mu_t < r_t$. Letting $n_t = N_t$, $\pi_t = q_t \delta_t / N_t$, $\mu_t - r_t = E_t(dr_t^D - dr_t)$, and $\sigma_t = (\sigma + \sigma_t^q)$ gives the desired result.

C.6 Proof of Proposition 6

As in the unconstrained case, the equilibrium features two regimes. First, I show that whenever it is optimal for experts to fully finance the risky projects, $\psi_t = 1$, then the VaR constraint cannot be binding. To see why, notice that whenever $\psi_t = 1$, market clearing still requires that $q_t = \bar{q} = (a + 2\theta)/\rho$, so $\mu_t^q = \sigma_t^q = 0$. From Proposition 6, Expert optimality in this case means that

$$\frac{\psi_t}{\eta_t} = \min \left\{ \frac{a/q_t + g + \mu_t^q + \sigma\sigma_t^q - r_t}{(\sigma + \sigma_t)^2}, \frac{1}{\alpha(\sigma + \sigma_t^q)} \right\}.$$

Since prices are constant, this condition requires that

$$\frac{1}{\eta_t} = \min \left\{ \frac{\rho + g - r_t}{\sigma^2}, \frac{1}{\alpha\sigma} \right\}.$$

If the constraint binds, then $\eta_t = \alpha\sigma$. But the law of motion for η_t (which was derived without relying on expert optimality), is still given by

$$d\eta_t = \left[\frac{(1 - \eta_t)^2}{\eta_t} \sigma^2 + (a/\bar{q} - \rho)\eta_t + \theta/\bar{q} \right] dt + (1 - \eta_t)\sigma dZ_t.$$

Therefore, in order for $\eta_t = \alpha\sigma$ to hold, it must be that $\sigma_t^\eta = 0$, $\Rightarrow \alpha\sigma = 1$, which is a contradiction, by assumption. Therefore, the constraint must not bind. This is an intuitive result: if the experts are willing to finance all of the projects in the economy, the Lagrange multiplier on the VaR constraint should not be positive. Therefore, the equilibrium when $\psi_t = 1$ is exactly the same as before.

In the other regime, however, things may be different. Fortunately, the equilibrium is still technically easy to solve, on account of the logarithmic utility. As before, note that when $\psi_t \in (0, 1)$, both the household and expert optimality conditions must hold with equality.

Since household portfolio optimality is unchanged,

$$\frac{1 - \psi_t}{1 - \eta_t} = \frac{\underline{a}/q_t + \underline{g} + \mu_t^q + \sigma\sigma_t^q - r_t}{(\sigma + \sigma_t)^2}.$$

Combining this with expert optimality using Proposition 6, in order to remove the $\mu_t^q + \sigma\sigma_t^q - r_t$ term, implies that

$$\frac{\psi_t}{\eta_t} = \min \left\{ \frac{(a - \underline{a})/q_t + (g - \underline{g})}{(\sigma + \sigma_t^q)^2} + \frac{1 - \psi_t}{1 - \eta_t}, \frac{1}{\alpha(\sigma + \sigma_t^q)} \right\}.$$

Lastly, notice that all of the other conditions which follow from Itô's Lemma still apply exactly as before.

C.7 Numerical methods

I solve the model using a suite of programs built off of the resources provided along with [Sannikov \(2013\)](#). Here I outline the basics of the procedure. The normal regime in both cases is available in closed form. The crisis regime, however requires solving one ordinary differential equation (ODE). First, consider the unconstrained case. From Proposition 4, notice that in the crisis regime, the function $q'(\eta)$ can be evaluated given q and η using only the first three conditions

$$\begin{aligned} \psi(\eta) &= \frac{\rho q(\eta) - \underline{a} - 2\theta}{a - \underline{a}} \\ \sigma^q(\eta) &= \left[\frac{(a - \underline{a})/q(\eta) + (g - \underline{g})}{\psi(\eta)/\eta - (1 - \psi(\eta))/(1 - \eta)} \right]^{1/2} - \sigma \\ q'(\eta) &= \left(1 - \frac{\sigma}{\sigma + \sigma^q(\eta)} \right) \frac{q(\eta)}{\psi(\eta) - \eta}. \end{aligned}$$

Therefore, I can use MATLAB's built-in function `ode45.m` to solve this ODE forward starting from initial condition $q(\epsilon) = (\underline{a} + 2\theta)/\rho$, where $\epsilon > 0$ is a small value chosen to avoid the singu-

larity that arises when $\eta = 0$. I impose the value-matching condition that $\max_{\eta \in [\epsilon, 1]} q(\eta) = \bar{q}$ to ensure that this regime coincides with the second one at $\eta = \eta^*$. Once I have found the function $q(\eta)$ in vector form, I can use numerical differentiation to compute $q''(\eta)$ and recover μ_t^q and r_t .

In the constrained model, the situation is exactly the same, except that σ_t^q can not be solved for explicitly as before.

$$\begin{aligned}\psi(\eta) &= \frac{\rho q(\eta) - \underline{a} - 2\theta}{a - \underline{a}} \\ \psi(\eta) &= \min \left\{ \frac{[(a - \underline{a})/q(\eta) + (g - \underline{g})]\eta}{(\sigma + \sigma^q(\eta))^2} + \frac{1 - \psi(\eta)}{1 - \eta}, \frac{1}{\alpha(\sigma + \sigma^q(\eta))} \right\} \\ q'(\eta) &= \left(1 - \frac{\sigma}{\sigma + \sigma^q(\eta)} \right) \frac{q(\eta)}{\psi(\eta) - \eta}\end{aligned}$$

To overcome this difficulty, I implement the following algorithm to compute $q''(\eta)$ given q and η :

1. Initialize (q, η) and set $\psi = (\rho q - \underline{a} - 2\theta)/(a - \underline{a})$.
2. Define the variables:

$$\begin{aligned}x_1 &\equiv \left[\frac{(a - \underline{a})/q + (g - \underline{g})}{(\psi/\eta - (1 - \psi)/(1 - \eta))} \right]^{1/2} \\ x_2 &\equiv \left[\frac{1}{\alpha\psi/\eta} \right].\end{aligned}$$

3. If the following equality holds:

$$\frac{\psi}{\eta} = \min \left\{ \frac{(a - \underline{a})/q + (g - \underline{g})}{x_1^2} + \frac{1 - \psi}{1 - \eta}, \frac{1}{\alpha x_1} \right\}$$

then set $q'(\eta) = (1 - \sigma/x_1)q/(\psi - \eta)$, otherwise $q'(\eta) = (1 - \sigma/x_2)q/(\psi - \eta)$.

This allows me to compute $q(\eta)$ numerically, again using `ode45.m`. I can then use numerical differentiation to recover $q''(\eta)$ as before.

Once I have obtained numerical values for the parameters $\sigma^\eta(\eta_t)$ and $\mu^\eta(\eta_t)$, I can simulate time-series paths for $\{\eta_t\}$ by discretizing the law of motion for η_t using the Euler-Maruyama method. Specifically, I want to simulate

$$d\eta_t = \mu^\eta(\eta_t)dt + \sigma^\eta(\eta_t)dZ_t.$$

To do this, I choose a grid $[0, T]$ and increment size dt (in my simulations, I choose $\eta_0 = 0.5$, $T = 5000$, and $dt = 1/12$), and use the following algorithm:

1. Initialize η_0 ,
2. Draw $\varepsilon \sim \mathcal{N}(0, 1)$ and set

$$\eta_1 = \eta_0 + \mu^\eta(\eta_0)dt + \sigma^\eta(\eta_0)\sqrt{dt}\varepsilon.$$

Since $\mu^\eta(\eta)$ and $\sigma^\eta(\eta)$ are defined on a grid, I use cubic spline interpolation to evaluate it for any value of $\eta \in [0, 1]$.

3. Repeat step 2 to given η_1 to obtain η_2, η_3, \dots

I use a burn-in period of 1000×12 , to eliminate dependence on the initial condition; furthermore, to ensure that my simulated values are in fact draws from the stationary distribution, I repeat this exercise 1000 times and experiment with different starting values η_0 . The distributions I plot are then kernel density estimates of the simulated draws.